SPERRY UNIVAC

1100 Series

8405/8430/8433 Disc Subsystem

Programmer Reference

SPERRY UNIVAC COMPUTER SYSTEMS

UP–8324

This document contains the latest information available at the time of publication. However, Sperry Univac reserves the right to modify or revise its contents. To ensure that you have the most recent information, contact your local Sperry Univac representative.

Sperry Univac is a division of Sperry Rand Corporation.

PRINTED IN U.S.A.

PAGE STATUS SUMMARY

ISSUE: UP–8324

Part/Section	Page Number	Update Level	Part/Section	Page Number	Update Level	Part/Section	Page Number	Update Level
Cover/Disclaimer								
PSS	1							
Contents	1 thru 5							
1	1 thru 3							
2	1 thru 24							
3	1 thru 41							
Appendix A	1 thru 13							
User Comment Sheet								
Total: 89 pages and covers								

Contents

FIGURES

TABLES

1. Introduction

1.1. GENERAL

This manual contains programming information for the SPERRY UNIVAC 8405/8430/8433 Disc Subsystem for SPERRY UNIVAC 1100 Series Systems. It is assumed that the programmer has sufficient background on the 1100 Series Systems and requires instruction only in the use of the 8405/8430/8433 Disc Subsystem. Operating information for the 8405/8430/8433 Disc Subsystem is included in the SPERRY UNIVAC 8405/8430/8433 Disc Subsystem Operator Reference, UP–8323 (current version).

1.2. SPERRY UNIVAC 8405/8430/8433 DISC SUBSYSTEM

The SPERRY UNIVAC 8405/8430/8433 Disc Subsystem (disc subsystem) consists of a 5039 Storage Control Unit (SCU) and from one to eight 8405 Fixed Head Discs (8405 FHD), an SCU and from two to sixteen 8430 Disc Units (8430 disc) or an SCU and a mix of one to eight 8405 FHDs and two to eight 8430 discs and/or 8433 discs. The SCU is shown in Figure 1–1, the 8405 FHD is shown in Figure 1–2, and the 8430 and 8433 discs are shown in Figure 1–3. The disc subsystem provides word-oriented processors with random or sequential access to large data files for online use. Compatibility between the 8-bit byte-oriented disc subsystem and the word-oriented processor is provided by two versions (buffered or unbuffered) of the SCU which includes the provision for adapting the 8-bit byte-oriented disc subsystem to the 36-bit word-oriented 1100 Series Processors.

1.2.1. Disc Units

The 8405 FHD contains a single fixed-head disc drive and a nonremovable disc stack. The 8430 disc and 8433 disc each contain a single disc drive with removable disc packs.

1.2.2. Subsystem Mixes

The 8405 FHD, 8430 disc, and the 8433 disc may be mixed in the same disc subsystem. The mix may be:

- two to sixteen 8430 discs and/or 8433 discs,
- one to eight 8405 FHDs, or
- one to eight 8405 FHDs and two to eight 8430 discs and/or 8433 discs.

1.2.3. Disc Subsystem Characteristics

The 8405 FHD, 8430 disc, and 8433 disc:

- receive control signals and data from the SCU,
- notify the SCU of operating conditions within a disc, and
- receive and send data read from a disc pack (or disc stack) to the SCU.

The SCU:

- receives and interprets command information from the processor,
- accepts data from an addressed disc drive with its associated disc pack, and makes data available to the processor,
- requests data from a processor and controls the writing of that data on the disc pack (or disc stack) in the addressed disc unit,
- detects and corrects data errors,
- adapts the 8-bit byte-oriented disc subsystem to the 36-bit word-oriented processors,
- provides Angular Position Sensing (APS),
- provides Command Retry, and
- provides 1024 words of buffer storage to simplify and improve data transfers with the processor or I/O unit.

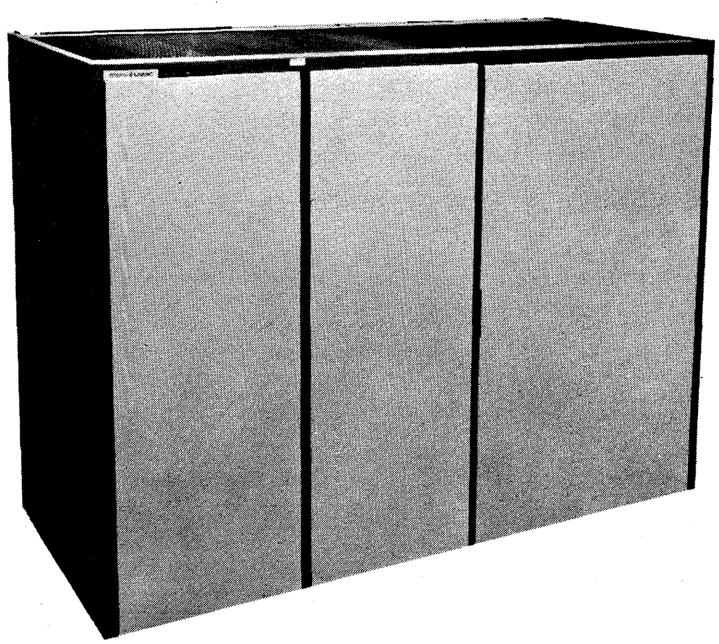

Figure 1–1. SPERRY UNIVAC 5039 Storage Control Unit

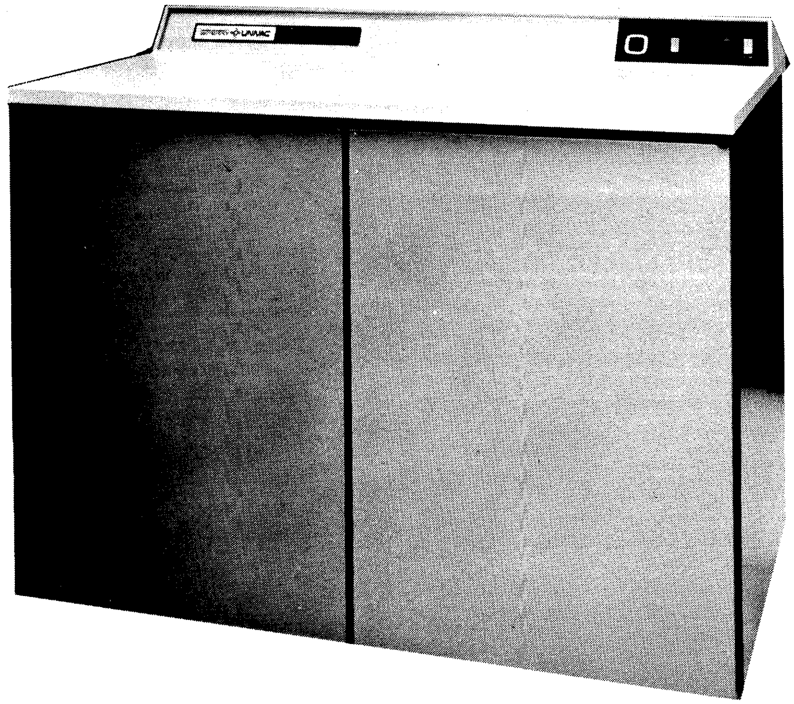

Figure 1–2. SPERRY UNIVAC 8405 Fixed Head Disc

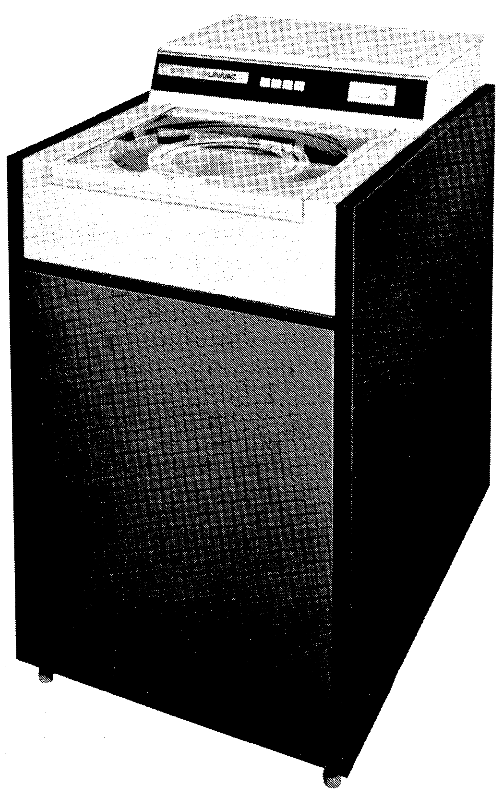

Figure 1–3. SPERRY UNIVAC 8430/8433 Disc Unit

2. Subsystem Description

2.1. GENERAL

The SPERRY UNIVAC 8405/8430/8433 Disc Subsystem (disc subsystem) is a sequential and random access storage subsystem with disc units using removable disc packs (8430/8433 disc subsystem) or a nonremovable disc stack (8405 fixed head disc subsystem). The 8430/8433 disc subsystem removable disc packs are interchangeable with disc packs of other 8430/8433 disc subsystems, respectively. However, disc packs are not interchangeable between the 8430 disc subsystem and the 8433 disc subsystem. The 5039 Storage Control Unit (SCU) provides data conversion, instruction control for reading data from and writing data on disc surfaces of the disc packs, and converts the byte-oriented disc subsystem information into words for use with word-oriented processors and from words to bytes for return information.

2.2. CHARACTERISTICS

2.2.1. 8405 Fixed Head Disc Subsystem

The 8405 fixed head disc subsystem provides processors with the capability of up to 49.5 million (8405–00/01) or 24.7 million (8405–04/05) 8-bit bytes of online storage. The average access time to any one record is 8.4 milliseconds. Once accessed, data is transferred at a nominal rate of 622,000 bytes per second. The performance characteristics of the 8405 disc subsystem are summarized in Table 2–1.

2.2.2. 8430/8433 Disc Subsystem

The 8430/8433 disc subsystem provides processors with the capability of 1.238 billion 8-bit bytes (8430 disc) or 2.47 billion 8-bit bytes (8433 disc) of online storage. The average access time to any one record is 35.3 milliseconds (average seek time + average rotational latency). Once accessed, data is transferred at a nominal rate of 806,000 bytes per second. The performance characteristics of the 8430/8433 disc subsystem are summarized in Table 2–2.

Table 2–1. SPERRY UNIVAC 8405 Fixed Head Disc Subsystem Characteristics

Parameter	Characteristics	
	8405–00/01 FHD	8405–04/05 FHD
Disc units per SCU	1 to 8	1 to 8
SCUs per disc subsystem	1 (single access), 2 (dual access)	1 (single access), 2 (dual access)
Disc drives per cabinet	1	1
Data read write heads per disc drive	864	432
Cylinders per disc stack	71	35
Head address per disc stack	11	11
Tracks per disc stack	864	432
Data bytes per disc stack	6,193,152	3,096,576
Access time (maximum)	16.7 milliseconds	16.7 milliseconds
Access time (average)	8.4 milliseconds	8.4 milliseconds
Rotational latency time (average)	8.4 milliseconds	8.4 milliseconds
Data transfer rate	622,000 bytes per second 138,222 36-bit words per second	622,000 bytes per second 138,122 36-bit words per second
Track capacity (112 word records)	8,064 8-bit bytes 1,792 36-bit words	8,064 8-bit bytes 1,792 36-bit words
Storage capacity (per unit)	6,193,152 8-bit bytes 1,376,276 36-bit words	3,096,576 8-bit bytes 688,128 36-bit words
Disc subsystem capacity (8 units)	49,545,216 8-bit bytes 11,010,048 36-bit words	24,772,608 8-bit bytes 5,505,024 36-bit words
Angular positioning sector	128	128
Data bytes per angular positioning sector	63	63
Parity (over interface lines)	Odd byte parity on command, address, status, and data. Error correction codes on all record files.	
Disc rotational speed	3600 RPM	3600 RPM

Table 2–2. SPERRY UNIVAC 8430/8433 Disc Subsystem Characteristics

Parameter	Characteristics	
	8430 Disc	**8433 Disc**
Disc units per SCUs	2 to 16	2 to 16
SCUs per disc subsystem	1 (single access), 2 (dual access)	1 (single access), 2 (dual access)
Disc drives per cabinet	1	1
Data read/write heads per disc drive	19	19
Data recording surfaces per disc pack	19	19
Tracks per disc surface	411	815
Cylinders per disc pack	411	815
Tracks per disc pack	7809	15,485
Access time (one cylinder)	7 milliseconds	10 milliseconds
Access time (average)	27 milliseconds	30 milliseconds
Access time (maximum)	50 milliseconds	55 milliseconds
Rotational latency time (average)	8.3 milliseconds	8.3 milliseconds
Rotational latency time (maximum)	16.7 milliseconds	16.7 milliseconds
Data transfer rate	806,000 bytes per second 179,111 36-bit words per second	806,000 bytes per second 179,111 36-bit words per second
Track capacity	20 records of 504 bytes 2,240 36-bit words	20 records of 504 bytes 2,240 36-bit words
Storage capacity (per disc pack)	77,374,084 8-bit bytes 17,194,240 36-bit words	154,748,168 8-bit bytes 34,388,480 36-bit words
Disc subsystem capacity (16 units)	1,237,985,344 8-bit bytes 275,107,840 36-bit words	2,475,970,688 8-bit bytes 550,215,680 36-bit words
Angular positioning sectors per disc track	128	128
Data bytes per angular positioning sector	78.75	78.75
Parity (over interface lines)	Odd byte parity on command, address, status, sense, and data. Error correction codes on all record files.	
Disc rotational speed	3600 RPM	3600 RPM

2.3. CONFIGURATION

The 8405/8430/8433 disc subsystem minimum operating configuration (Figure 2–1) consists of two 8430/8433 discs or one 8405 FHD and one SCU (buffered or unbuffered).

NOTE:

The minimum 8430/8433 disc configuration must consist of two 8430s or two 8433s, not one 8430 and one 8433. The minimum configuration can be expanded to include up to:

- *sixteen 8430/8433 discs per SCU, or*
- *eight 8405 FHDs per SCU, or*
- *eight 8405 FHDs and eight 8430/8433 discs per SCU.*

A dual-access disc subsystem (Figure 2–2) requires two SCUs and two disc units equipped with the dual-access capability. The dual-access configuration allows simultaneous read/read, read/write, write/read, or write/write operations on two disc units. Table 2–3 is a summary of disc subsystem components and features.

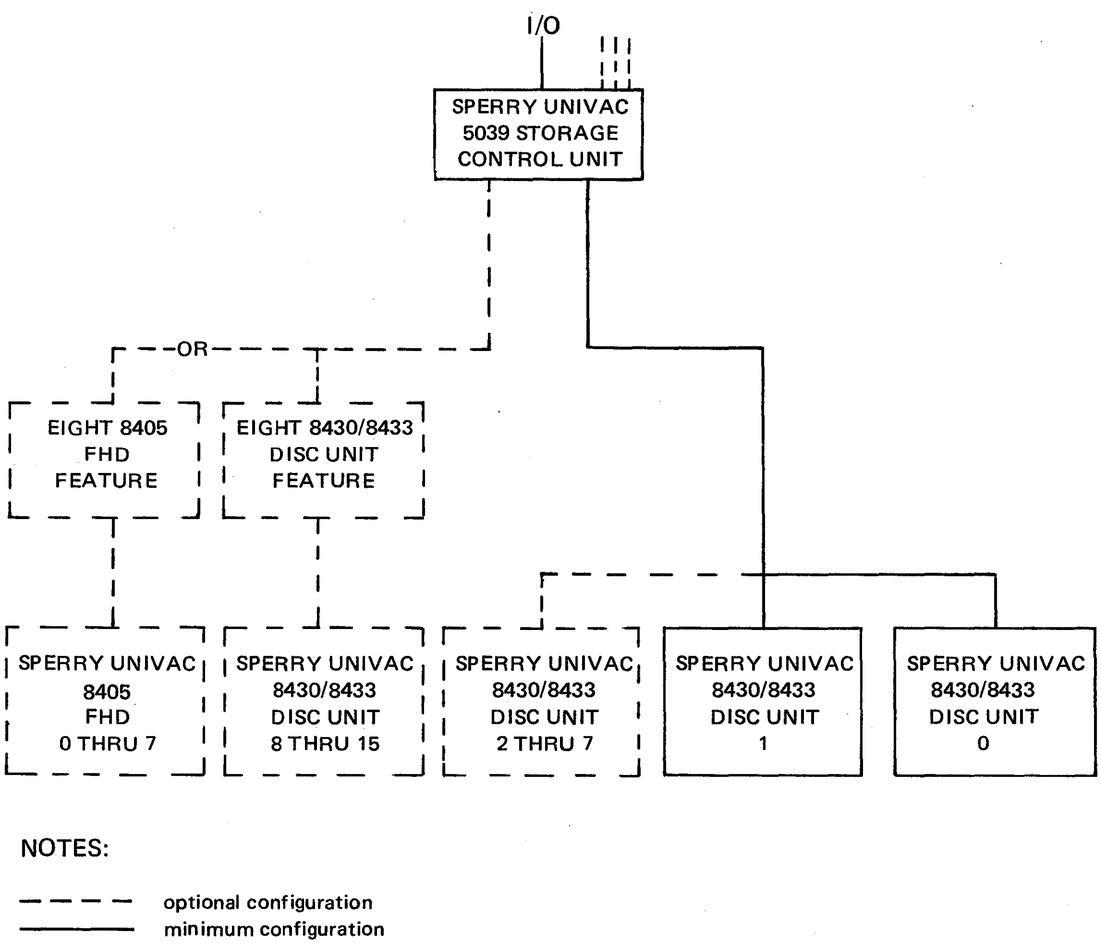

Figure 2–1. SPERRY UNIVAC 8405/8430/8433 Disc Subsystem Configuration

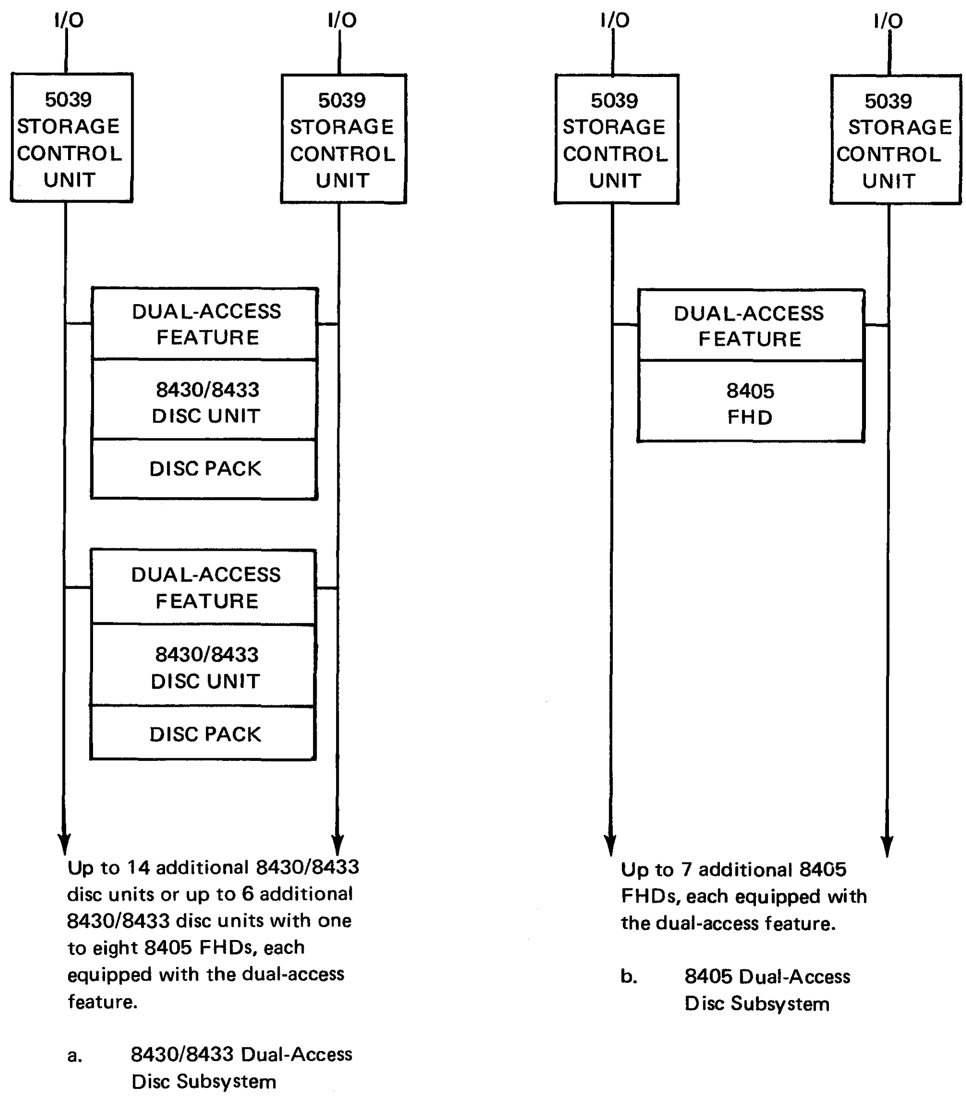

Figure 2–2. SPERRY UNIVAC 8405/8430/8433 Disc Subsystem with Dual Access Capabilities

Table 2–3. SPERRY UNIVAC 8405/8430/8433 Disc Subsystem Components and Features

Name	Description
5039 Storage Control Unit	Controls up to eight 8430/8433 discs and, with the eight 8405 disc feature, up to eight 8405 FHDs. With the eight 8430/8433 disc feature, up to sixteen 8430/8433 discs can be controlled.
8405 FHD	Provides a single fixed head disc drive with a nonremovable disc stack.
8430 Disc Unit	Provides a single disc drive using a removable disc pack.
8430 Disc Pack	Provides up to 77 million bytes of removable storage.
8433 Disc Unit	Provides a single disc drive using a removable disc pack.
8433 Disc Pack	Provides up to 154 million bytes of removable storage.
Dual Access	Provides dual access and simultaneous read/read, read/write, write/read, or write/write operations on any two discs. Requires two SCUs.
SPI	Provides nonsimultaneous access to the SCU from two selector channels and, also adds third and fourth channel to dual channel feature for four channel capability.
16 Drive Expansion	Provides for attaching up to sixteen 8430/8433 disc drives to the SCU. This feature is mutually exclusive of the 8405 capability feature.
8405 Capability	Adds the capability to allow attachment of from 1 to 8 additional 8405 FHDs to the SCU. This feature is mutually exclusive of the 16 Drive Expansion feature.

2.4. SUBSYSTEM COMPONENTS

A subsystem contains at least one or two (depending on type) disc units, one storage control unit, and one disc pack for each 8430/8433 disc. (The 8405 disc subsystem requires only one disc unit; the 8430/8433 disc subsystems require two disc units, either two 8430 discs or two 8433 discs.)

Each disc unit contains a disc drive and disc packs (or disc stack). The disc drive is the mechanical drive for a disc pack and the disc pack is the storage medium in the form of magnetic discs. The storage control unit provides the electronic control for writing and retrieving data from a particular disc unit and disc pack; and converts the disc subsystem bytes information to the word format required by the word-oriented processor.

2.4.1. Disc Units

The disc units operate as slaves to the 5039 Storage Control Unit (SCU). The disc units store (write) information onto or retrieve (read) information from the disc packs. The disc unit contains the physical set of electronics required to:

- rotate the disc pack (8430/8433) or disc stack (8405) at the proper speed,
- position the read/write heads (perform a seek) at a particular cylinder (moving-head disc units, 8430/8433 discs),
- select the read/write heads (8405 disc),

- transfer data from/to the accessed area on the disc pack (read/write),
- monitor disc unit operation, generate status information, and transfer information on the status of an operation to the storage unit, and
- buffer data transfers using integral buffer storage.

The disc unit, in order to seek or transfer data, must be addressed by the processor.

2.4.1.1. 8405 Fixed Head Disc

The 8405 Fixed Head Disc (FHD) is a random access, fixed-head disc file with one read/write head for each track. The 8405–00/01 disc has 864 addressable tracks on 12 recording surfaces. Heads are selected by addressing cylinders 00 through 63 and heads 00 through 11. The 8405–00/01 disc capacity is 6,193,152 bytes. Data is transferred to and from the SCU at 622K byte/sec. The 8405–04/05 disc has 432 addressable tracks on 6 recording surfaces. The 8405–04/05 disc capacity is 3,096,576 bytes. The data transfer rate is the same as that of 8405–00/01 disc.

2.4.1.2. 8430 Disc Unit

The 8430 disc unit (8430 disc) is a random access, moving-head disc unit featuring a removable and interchangeable disc pack (F1230–00) with a gross capacity of 100 megabytes (one record per track) or a net capacity of 77.3 megabytes (20 records per track). The disc pack is not compatible with other disc types. Figure 2–3 shows the disc pack cylinder arrangement. Nineteen disc pack data surfaces are serviced by one moving actuator mechanism which contains one read/write head for each surface. Each data surface contains 411 data tracks. The moving actuator mechanism utilizes information from the servo surface to position all nineteen read/write heads at one of the 411 cylinders. The time required to move from one nineteen-track cylinder to another is seven milliseconds (to the adjacent cylinder), 27 milliseconds (average), and 50 milliseconds (full travel). Average rotational latency is 8.33 milliseconds. Angular position sensing information is supplied to the SCU to allow channel release while moving the actuator and searching for a particular sector on a track. Data is transferred to and from the disc pack at a 6.45 MHz bit rate.

The dual-access feature allows the physical attachment of two SCUs to an 8430 disc. This dual-access feature allows an active 8430 disc to be interrogated as to status and availability. A priority circuit in the drive provides a tie-breaking function should two SCUs attempt to select the same drive at the same time.

2.4.1.3. 8433 Disc Unit

The 8433 disc unit (8433 disc) is essentially a double track density 8430 disc. It uses the F1223–00 disc pack, which is removable and interchangeable between 8433 discs. It is not compatible with other disc types. Figure 2–4 shows the disc pack cylinder arrangement. Disc pack gross capacity is 200 megabytes (one record per track) or a net capacity of 154.7 megabytes (20 records per track). Read/write head positioning time is the same as for the 8430 disc except that the average access time to a cylinder is 30 milliseconds. The F1223–00 disc pack contains 815 cylinders of 19 tracks each. The dual-access feature for the 8433 disc operates essentially the same as the 8430 disc dual-access feature. The data transfer rate is the same as that of the 8430 disc.

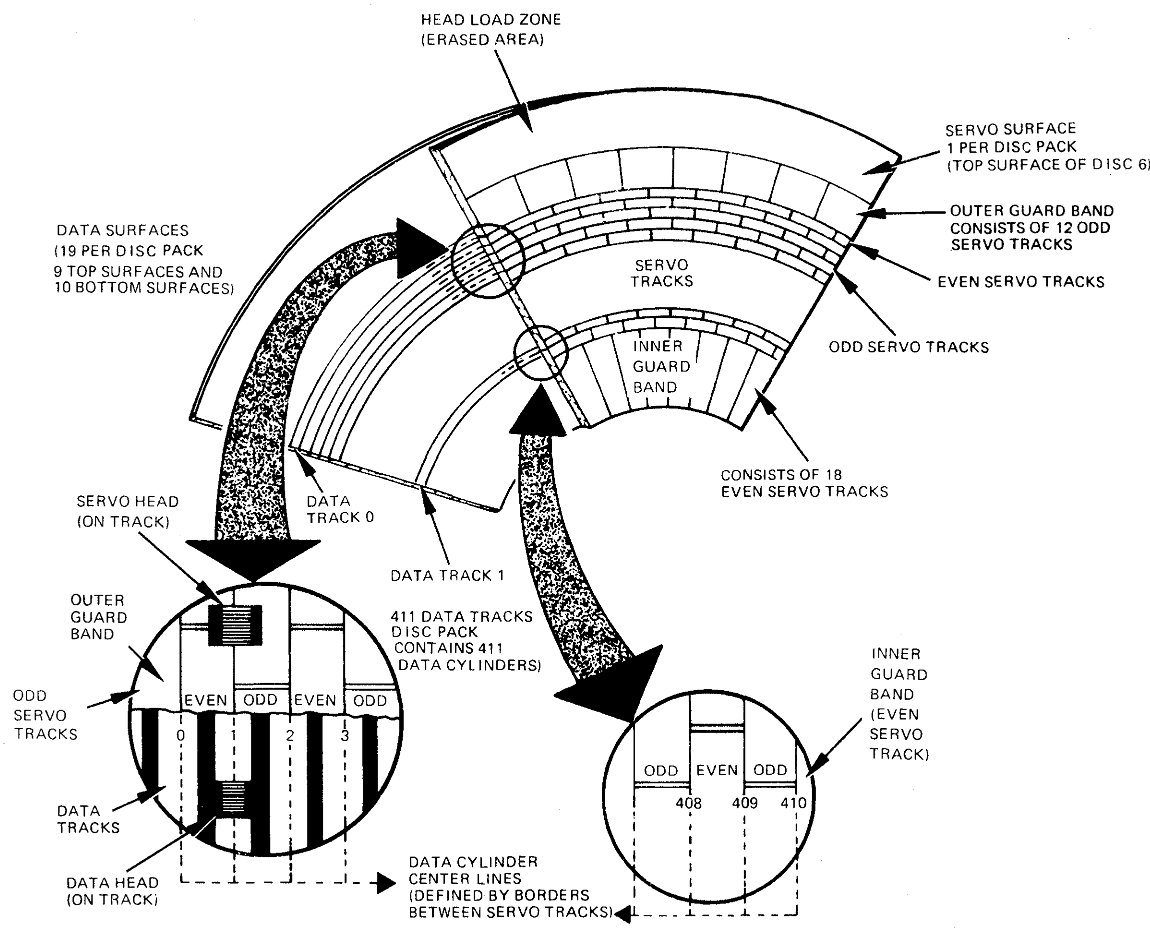

Figure 2–3. SPERRY UNIVAC 8430 Disc Pack (F1230–00) Disc Cylinder Assignment

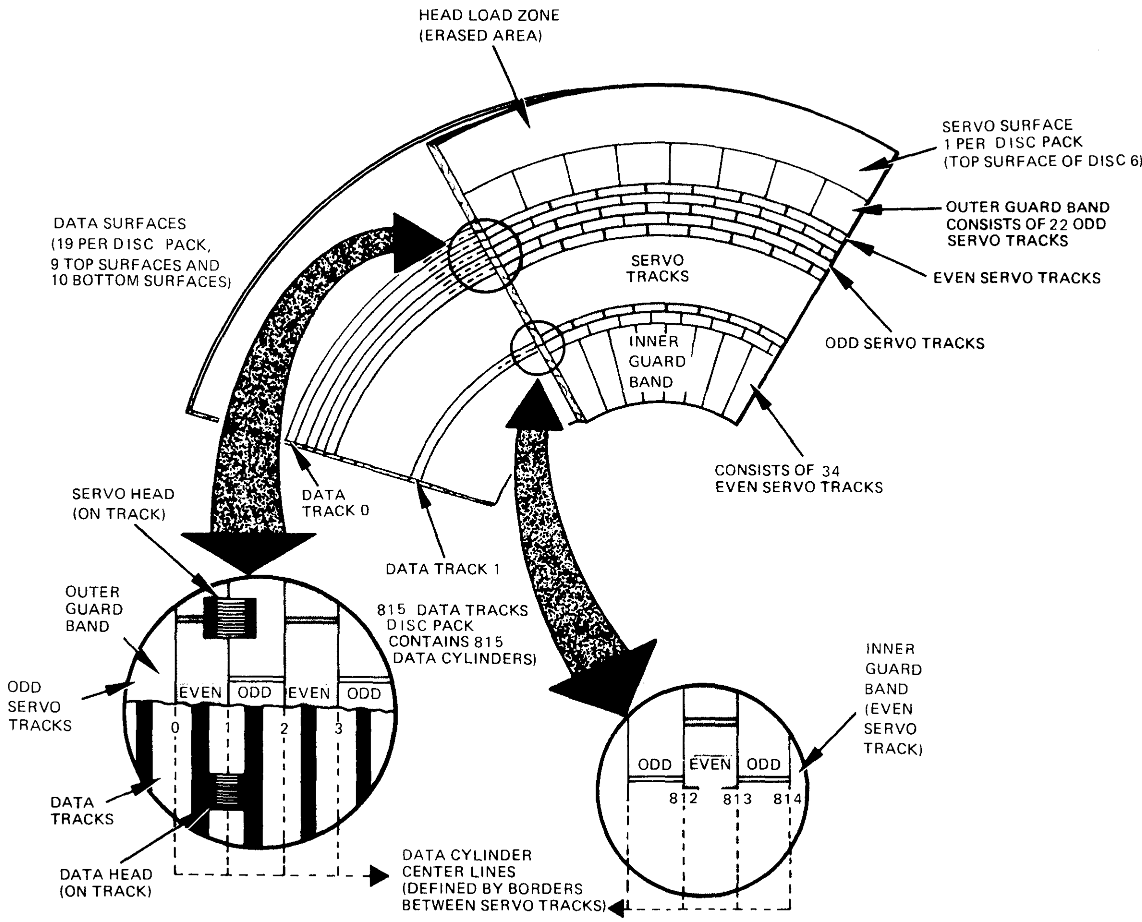

Figure 2–4. SPERRY UNIVAC 8433 Disc Pack (F1223–00) Disc Cylinder Assignment

2.4.1.4. Disc Unit Logical Addressing (Programmer Selectable)

The operator control panel, on the disc unit, includes a removable module select plug which establishes the logical address for the unit. These module select plugs are numbered 0–7 for up to eight disc units. With the sixteen disc unit feature added, the module select plugs are numbered 0–F (Hexadecimal). The module select plugs are removable and interchangeable, thus permitting one physical disc unit to assume different logical addresses. The 8405 FHD logical addressing must be either 0–7 or 8–F, not 0–F, (Hexadecimal) and the module select is fixed (nonremovable). The logical device address is transmitted in the form of an address byte from the processor. The format of the address bytes is illustrated below:

0	1	2	3	4	5	6	7	Bit position
C	C	C	C	A	D	D	D	Bit name
Address Byte								

where CCCC is any combination of binary bits specifying one of 8 SCUs and ADDD identifies the disc unit.

NOTE:

Without the sixteen disc unit feature, A is always zero (0).

2.4.1.5. Disc Unit Physical Addressing (Hard-Wired)

At the time the disc subsystem is installed, a hard-wired physical address is assigned to each disc unit by the Sperry Univac customer engineer. This address identifies to the system the physical position the disc unit occupies relative to the SCU. The physical address is transmitted to the system when a Sense command is performed by the SCU. The physical address is determined from a 3 of 6 code as shown below:

- Disc unit A – 111000
- Disc unit B – 110001
- Disc unit C – 101010
- Disc unit D – 100011
- Disc unit E – 011100
- Disc unit F – 010101
- Disc unit G – 001110
- Disc unit H – 000111

With the sixteen disc unit feature installed the SCU separately identifies which bank of eight disc units is being referenced and then utilizes the physical address to uniquely determine the individual disc unit.

2.4.1.6. Disc Unit Access Arm Assembly (8430 and 8433 only)

The 8430/8433 disc unit access arm assembly consists of an array of read/write heads used to transfer data and a positioning servo system to move the read/write heads to the desired location.

Positioning (Seeking) is controlled by use of the read only servo head which determines track locations from the servo head which in turn determines track locations from the servo surfaces on the disc pack (top surface of the sixth disc).

Once the desired cylinder and track have been reached (Seek Complete) the selected head is enabled to read or write.

2.4.2. 8430/8433 Disc Pack

The 8430 disc pack (F1230–00) and the 8433 disc pack (F1223–00) each consist of 12 discs on a vertical shaft, or spindle, with 19 disc surfaces for data recording and one for access positioning. The outside surfaces of the top and bottom discs are used to protect the inner surfaces and are not used for data recording. Circular protective plates also are mounted below the bottom disc and above the top disc to protect the disc pack from mechanical damage. A dust cover is an integral part of the removable disc pack cover. The disc pack is placed on the spindle and the handle locks the disc pack to the spindle before the cover and handle are removed. The entire assembly of discs, spindle and protective plates rotate at 3600 revolutions per minute. The disc packs are removable and interchangeable within same types. Figure 2–5 illustrates an 8430 disc pack and Figure 2–6 illustrates an 8433 disc pack.

CAUTION

Damage to the spindle or the disc pack is possible if an attempt is made to mount the wrong disc pack in a disc unit.

2.4.2.1. Disc Pack Organization and Addressing Format

Data on the disc pack must be accessed by use of cylinder and track addresses. An additional level of addressing (angular position sensing) is provided to increase the number of discs which can be serviced efficiently by one SCU.

2.4.2.1.1. Tracks and Cylinders

Each recording surface of the disc pack is assigned its own read/write head. Access to data on the disc pack is obtained by positioning the array of read/write heads at any one of the 411 (8430 disc) or 815 (8433 disc) disc positions. All read/write heads operate in the same vertical plane; therefore the tracks are numbered 0 to 18 from top to bottom defining one cylinder. Figure 2–7 is a cutaway of the entire 8430 disc pack and Figure 2–8 is a cutaway of the entire 8433 disc pack. The concentric cylinders are numbered 000 to 410 for the 8430 disc and 000 to 814 for the 8433 disc starting at the periphery and moving toward the center; therefore the address of an individual track in a given disc pack consists of a cylinder number and a read/write head (track) number. The addressing format requires that four 8-bit bytes be used to specify a complete cylinder and track address. This format is illustrated in Figure 2–7 for the 8430 disc and Figure 2–8 for the 8433 disc.

Access to different tracks within a cylinder is faster than access to tracks in different cylinders since changing tracks requires only read/write head electronic switching; whereas accessing a different cylinder requires physical movement of the access arm assembly.

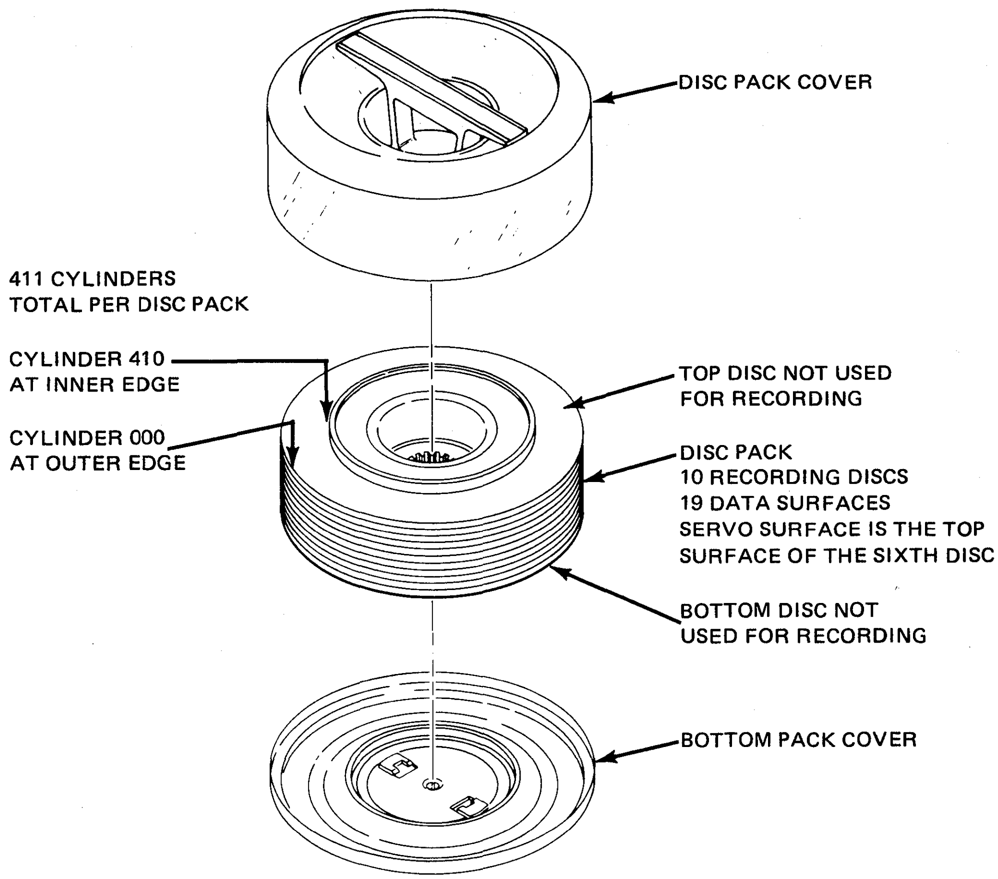

NOTE:

Each cylinder position intersects a disc data surface to establish a track location.

Figure 2–5. SPERRY UNIVAC 8430 Disc Subsystem Disc Pack (F1230–00)

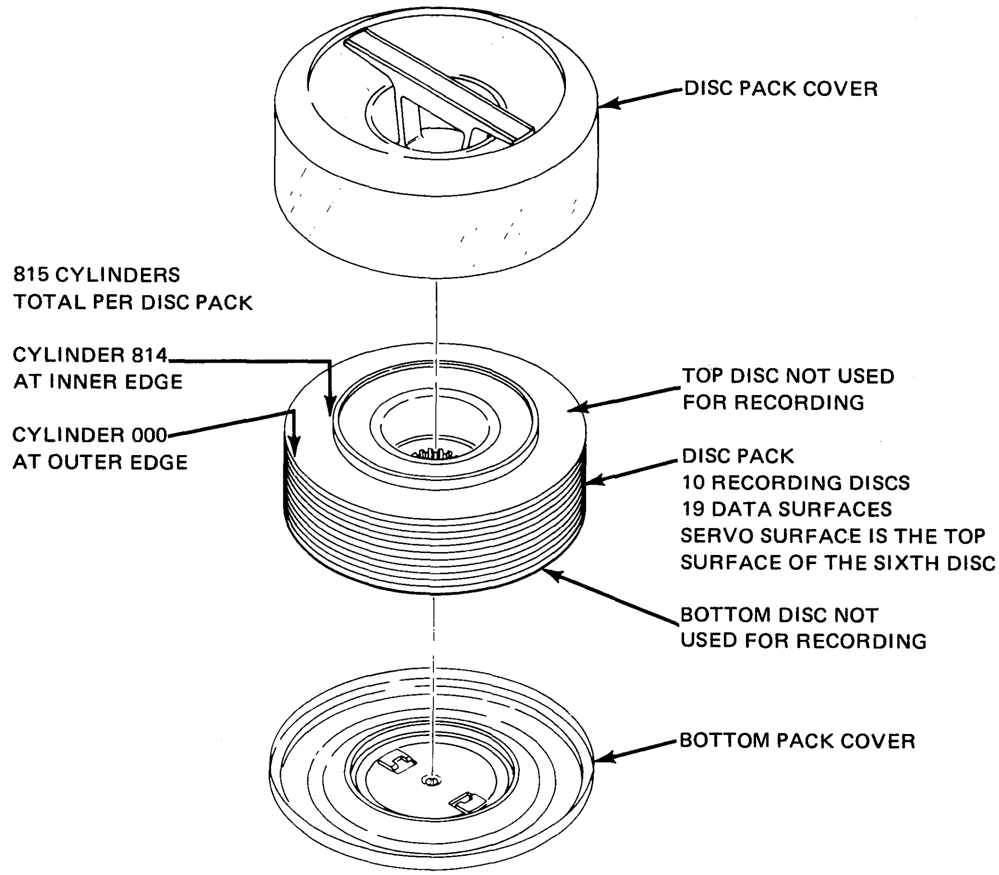

NOTE:

Each cylinder position intersects a disc data surface to establish a track location.

Figure 2–6. SPERRY UNIVAC 8433 Disc Subsystem Disc Pack (F1223–00)

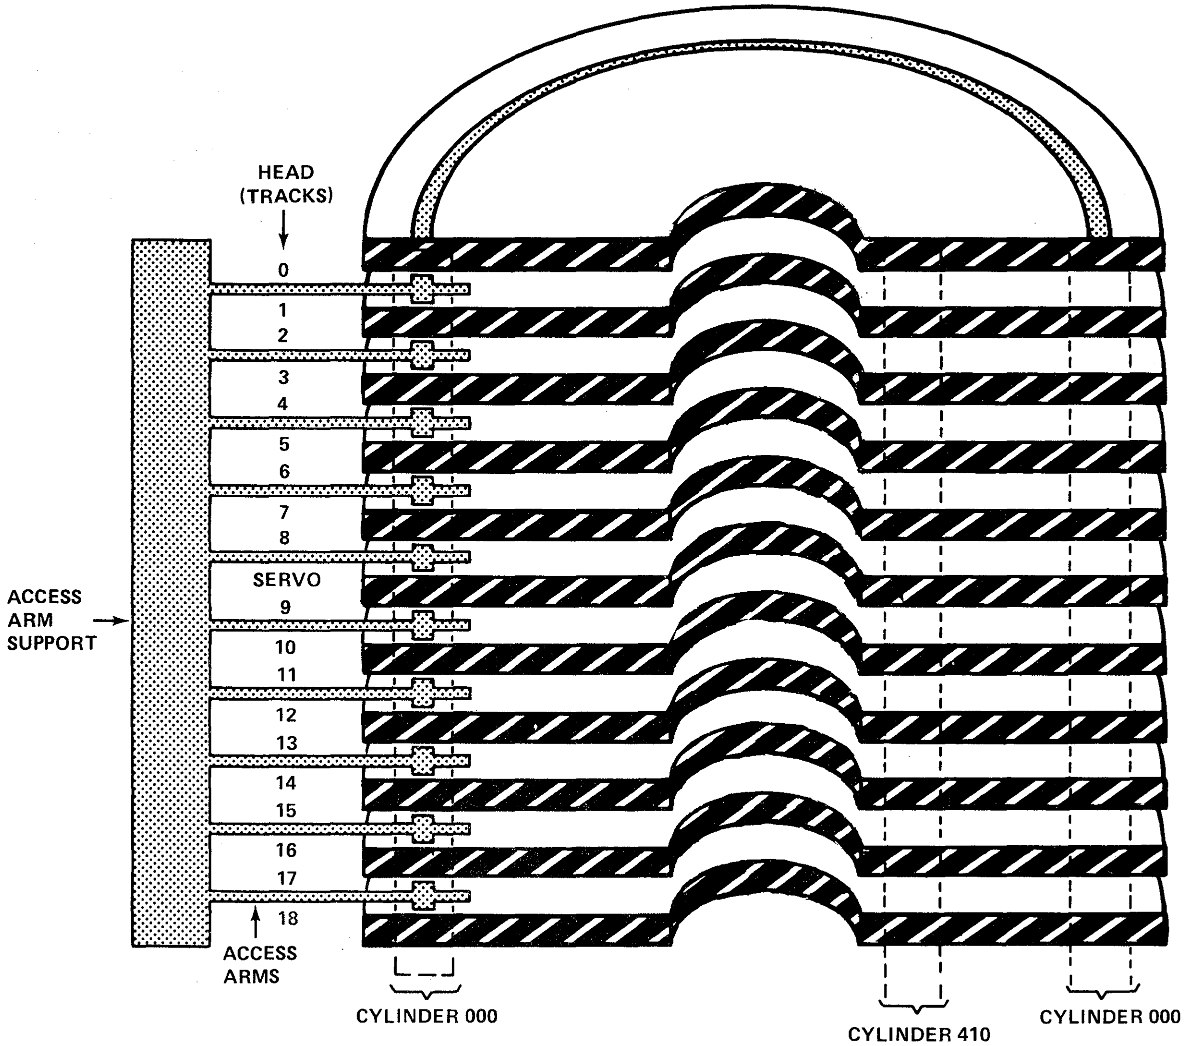

Bit position	0	1	2	3	4	5	6	7
Bit value	Not used							256
	High order cylinder address bits 0–6 must be zero							
	CYLINDER Byte 1							

Bit position	0	1	2	3	4	5	6	7
Bit value	128	64	32	16	8	4	2	1
	Low order cylinder address							
	CYLINDER Byte 2							

Bit position	0	1	2	3	4	5	6	7
Bit value	Not used							
	High order head address bits 0–7 must be zero							
	HEAD Byte 3							

Bit position	0	1	2	3	4	5	6	7
Bit value	Not used			16	8	4	2	1
	Low order head address bits 0–2 must be zero							
	HEAD Byte 4							

Figure 2–7. SPERRY UNIVAC 8430 Disc Unit, Cylinder and Track Addressing

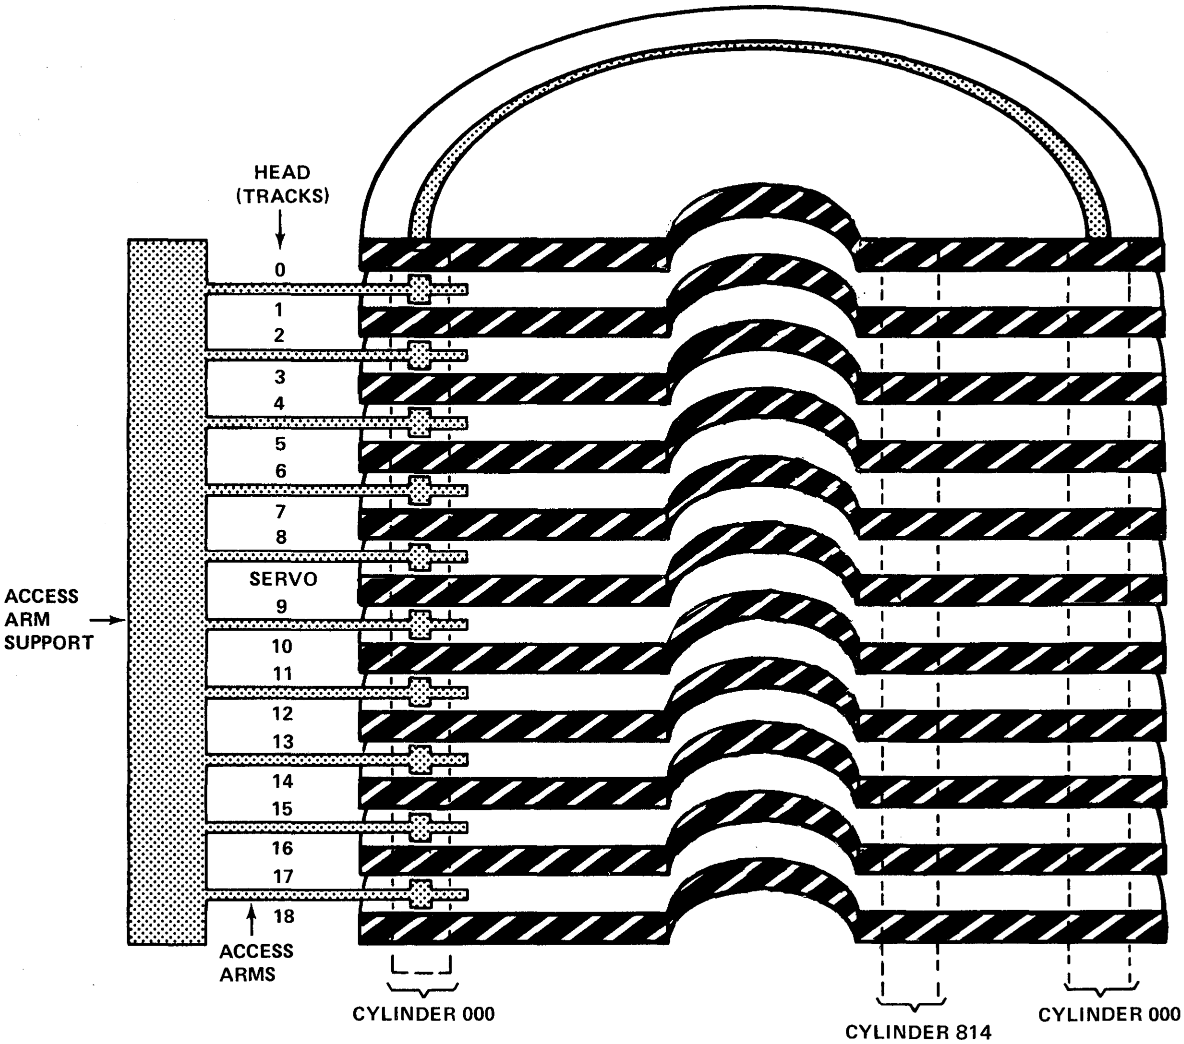

Bit position	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
Bit value	Not used						512	256	128	64	32	16	8	4	2	1	Not used								Not used			16	8	4	2	1
	High order cylinder address bits 0–5 must be zero								Low order cylinder address								High order head address bits 0–7 must be zero								Low order head address bits 0–2 must be zero							
	CYLINDER Byte 1								CYLINDER Byte 2								HEAD Byte 3								HEAD Byte 4							

Figure 2–8. SPERRY UNIVAC 8433 Disc Unit, Cylinder and Track Addressing

2.4.2.1.2. Angular Position Sensing

To achieve angular position sensing (APS), the disc unit electronically divides the 8430/8433 disc pack into 128 sectors numbered from 0 to 127 (Figure 2–9). On the 8430/8433 there are 64 odd sectors of 106 bytes each and 64 even sectors of 104 bytes each. The 8405 disc is divided into 128 odd and even sectors of 81 bytes each. The specified sector address is transmitted from the processor to the disc unit and when the sector is detected the processor is signaled. The sector address is transmitted in one byte and the format of that byte is illustrated in Figure 2–9. (See 3.2.1.4.2 and 3.2.1.5 for programming details).

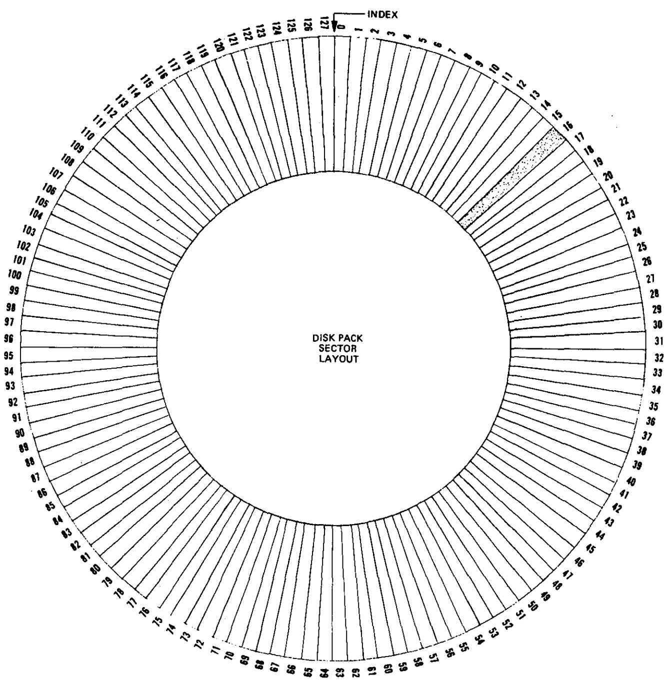

Bit position	0	1	2	3	4	5	6	7
Bit value	128	64	32	16	8	4	2	1
1 BYTE – SECTOR ADDRESS								

Figure 2–9. Angular Positioning Sectors

When the cylinder and head addressing is complete in the disc unit, the SCU would normally inform the processor the operation was complete. With angular position sensing activated, however, the addressing cycle is not complete until a specified angular position has been reached. Using angular position sensing, then, enables a record search to start in a programmer specified area of the track and thus reduce total record search time.

2.4.2.2. Disc Track and Record Formats

Each track consists of a home address, track descriptor record (RO), and one or more data records. The areas on the track are separated by gaps. One index mark serves the entire disc pack in indicating the physical beginning of the tracks. Figure 2–10 illustrates the complete track format from index to index.

(A) INDEX

A TRACK IS FROM INDEX POINT TO INDEX POINT (16.7MS) AND ALL TRACKS ON THE DISK PACK ARE SYNCHRONIZED BY THE SAME INDEX MARKER. INDEX MARKER INDICATES THE PHYSICAL BEGINNING OF EACH TRACK.

TRACK CAPACITY IS 13,440 BYTES AND THERE ARE 7,676 TRACKS PER PACK.

(B) GAPS 1,2,3 AND 4

GAPS ALLOW TIME FOR ONE COMMAND TO END AND THE CHANNEL TO TRANSFER A NEW COMMAND AND ALSO ALLOW TIME TO CHANGE FROM READ TO WRITE OR WRITE TO READ STATES. SYNC BYTES INDICATE THE BEGINNING OF DATA. ADDRESS MARKS START ALL RECORDS AFTER RO AND ARE USED TO FIND THE START OF A RECORD WITHOUT WAITING FOR INDEX

(C) DRIVE I.D. CODE (PHYSICAL ADDRESS) FOR CONTROL UNIT.

111000	DRIVE A
110001	DRIVE B
101010	DRIVE C
100011	DRIVE D
011100	DRIVE E
010101	DRIVE F
001110	DRIVE G
000111	DRIVE H

3 OF SIX DECODE. 3 BITS MUST BE ON TO BE VALID.

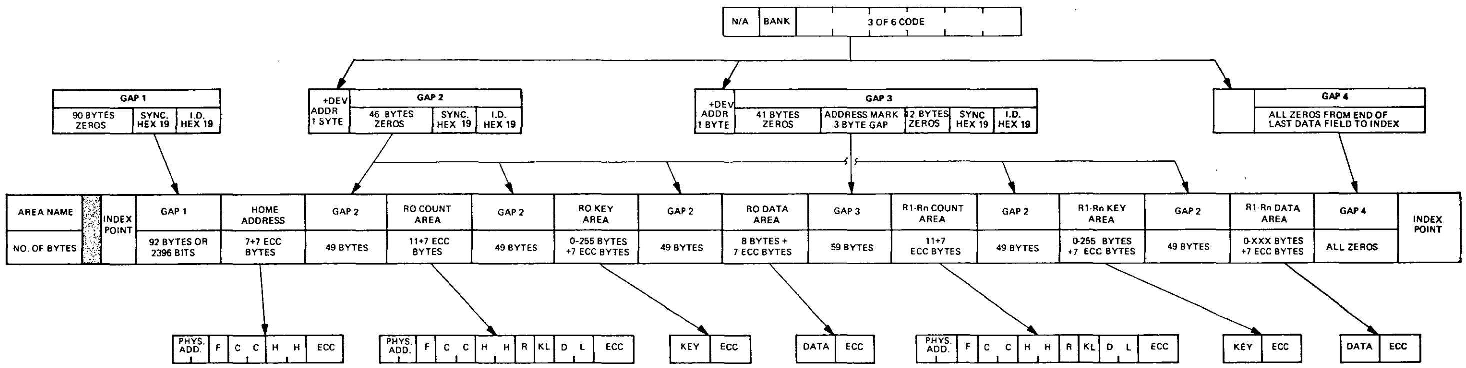

(D) HOME ADDRESS FIELD

PHYSICAL ADDRESS: TWO BYTES, PRECEDING HOME ADDRESS AND ALL COUNT FIELDS, INDICATE PHYSICAL ADDRESS AS FOLLOWS:

BYTE	BIT(S)	CONTENTS
0	0-7	LOW-ORDER CYLINDER ADDRESS
1	0	ZERO
1	1	HIGH-ORDER CYLINDER ADDRESS
1	2	ZERO
1	3-7	HEAD ADDRESS

F(FLAG): DEFINES CONDITION OF TRACK AND/OR INDICATES A CE DISC PACK. THIS IS THE ONLY FLAG BYTE TRANSFERRED TO OR FROM THE CHANNEL ON ALL OTHER FIELDS. THE C.U. WRITES THE FLAG USING THE FLAG READ FROM HOME ADDRESS OR OTHER COUNT FIELDS.
BITS 0–4 = ZERO
BIT 5 = 1,C.E. PACK
BITS 6 AND 7 - 00 = NORMAL TRACK
01 = ALTERNATE TRACK
10 = DEFECTIVE TRACK
11 = DEFECTIVE TRACK

CC(CYLINDER NUMBER): SPECIFIES CYLINDER NUMBER (FROM 0 TO 410)

HH (HEAD NUMBER): SPECIFIES THE READ/WRITE HEAD NO. WITHIN THE SELECTED CYLINDER (FROM 0 TO 18).

ECC(ERROR CORRECTION CODE): GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

GAPS: SEE GAP INFORMATION ABOVE.

*NOT USED BY IBM PROGRAMMING SYSTEMS.

(E) RECORD ZERO FIELD

RO COUNT AREA

PHYSICAL ADDRESS: SAME AS HOME ADDRESS FIELD

F(FLAG): DEFINES CONDITION OF THE TRACK AND INDICATES WHETHER OR NOT THIS IS AN OVERFLOW RECORD.
BITS 0 THROUGH 3–UNUSED AND WRITTEN AS ZEROS.
BIT 4–WHEN ON, INDICATES THAT THE RECORD CONTINUES ON NEXT TRACK
BIT 5–ALWAYS ZERO
BITS 6 AND 7–00 = NORMAL TRACK
01 = ALTERNATE TRACK
10 = DEFECTIVE TRACK

CC(CYLINDER NUMBER): SPECIFIES CYLINDER NUMBER (FROM 0 TO 410).

HH(HEAD NUMBER): SPECIFIES THE READ/WRITE HEAD NO. WITHIN THE SELECTED CYLINDER (FROM 0 TO 18).

R(RECORD NUMBER): SPECIFIES THE SEQUENTIAL NUMBER OF THE RECORD ON THE TRACK (ZERO IN THIS CASE)

KL(KEY LENGTH): SPECIFIES THE NUMBER OF BYTES IN THE RO KEY FIELD (FROM 0 TO 255 BYTES).

DL(DATA LENGTH): SPECIFIES THE NUMBER OF BYTES IN RO DATA FIELD (1 TO TRACK CAPACITY)

ECC(ERROR CORRECTION CODE): GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

* RO KEY AREA

KEYFIELD: IDENTIFIES INFORMATION IN THE DATA FIELD.

ECC(ERROR CORRECTION CODE): GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

RO DATA AREA

DATA FIELD: CONTAINS THE INFORMATION IDENTIFIED BY THE COUNT AND KEY AREAS. THIS VARIBLE DATA FIELD IS USED FOR CUSTOMER DATA AND CAN BE UP TO MAXIMUM FOR THE TRACK. IT IS ALSO USED FOR ALTERNATE OR DEFECTIVE TRACK ADDRESSES.

ECC(ERROR CORRECTION CODE): GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

GAPS: SEE GAP INFORMATION ABOVE.

(F) DATA RECORD FIELD

R1–Rn COUNT AREA

F(FLAG): SAME AS RECORD ZERO.

CC(CYLINDER NUMBER): SPECIFIES CYLINDER NUMBER (FROM 0 TO 410).

HH(HEAD NUMBER): SPECIFIES THE READ/WRITE HEAD NUMBER WITHIN THE SELECTED CYLINDER (FROM 0 TO 18).

R(RECORD NUMBER): SPECIFIES THE SEQUENTIAL NUMBER OF THE RECORD ON THE TRACK.

KL(KEY LENGTH): SPECIFIES THE NUMBER OF BYTES IN THE KEY FIELD (FROM 0 TO 255 BYTES).

DL(DATA LENGTH): SPECIFIES THE NUMBER OF BYTES IN THE DATA FIELD (FROM 1 TO TRACK CAPACITY).

ECC(ERROR CORRECTION CODE) GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

R1–Rn KEY AREA

KEYFIELD: IDENTIFIES INFORMATION IN THE DATA FIELD.

ECC(ERROR CORRECTION CODE): GENERATED BY THE S.C.U.–USED FOR ERROR DETECTION AND CORRECTION.

R1–Rn DATA AREA

DATA FIELD: THIS IS THE NORMAL CUSTOMER DATA FIELD AND IT CONTAINS THE INFORMATION IDENTIFIED BY THE COUNT AND KEY AREAS. DATA FIELD OF ZERO INDICATES END-OF-FILE (LOGICAL GROUP OF RECORDS). KEY AND/OR DATA FIELDS CAN BE REWRITTEN WITHOUT CHANGING BALANCE OF THE TRACK, BUT WRITING HOME ADDRESS, RO OR ANY COUNT FIELD REQUIRES THE BALANCE OF THE TRACK TO BE REWRITTEN OR ERASED.

ECC (ERROR CORRECTION CODE: GENERATED BY THE S.C.U. - USED FOR ERROR DETECTION AND CORRECTION

GAPS: SEE GAP INFORMATION ABOVE.

Figure 2–10. Track Format

2.4.2.2.1. Gaps

There are four different types of gaps within the disc track format designated G1 through G4. The primary purpose of gaps is to allow time for error correction code (ECC) bytes to be checked and form channel command turn-around time.

- Gap 1 – This gap is written between index and home address and is 92 bytes long.
- Gap 2 – This gap is 49 bytes long and is written between home address and R0 count; between the count and key areas of each record; between key and data areas; or if the record has no key, between the count and data areas of each record.
- Gap 3 – This gap is written between the data area of one record and the count area of the following record and is 59 bytes long.
- Gap 4 – This gap is written from the end of the last data field to the index and must be 40 or more bytes in length.

2.4.2.2.2. Home Address Area

The home address (HA) area defines track condition and reflects the physical location of the track within the disc pack. There is one HA area per track. This area, created by a Write HA operation is the first area written on each track.

The format of home address is:

INDEX POINT		G1	PA	F	C	C	H	H	ECC
No. of Bytes	0	92	2	1	1	1	1	1	7

1. PA (Physical Address)

   The track physical address bytes are generated by the SCU and are read and used exclusively by the SCU for seek verification.

   The format of PA is:

Bit	Byte 0 8405/8430/8433	Byte 1 8405	Byte 1 8430	Byte 1 8433
0	↑	0	Reverse Bit	0
1		0	Cyl (256)	Cyl (512)
2	Low	0	High Order Diff	Cyl (256)
3	Order	↑	↑	↑
4	Cylinder			
5		Head	Head	Head
6				
7	↓	↓	↓	↓

2. F (Flag Byte)

   The flag byte is usually reserved for internal use by the SCU. The exceptions are read home address and write home address where the flag byte is the first byte transferred.

   The bit assignments are:

Bit	Definition
0–4	Unused, written as zeros
5	If set indicates a CE pack for diagnostics only.
6–7	00 Normal track 01 Alternate track 10 Defective track 11 Defective alternate track

3. CCHH (Cylinder and Head Number)

   The CCHH field should be formatted to agree with the cylinder and head address for the track. The home address is not used by the SCU for Seek verification.

4. ECC (Error Correction Code)

   Error correction code bytes are generated and written by the SCU when the home address is written. The ECC bytes are verified whenever the home address field is read, searched or clocked (read internally by the SCU but not transferred to the channel).

5. Home Address Field

   The home address field is defined to be the CCHH portion of the home address area.

2.4.2.2.3. Track Descriptor Record (TDR)

R0 (TDR) is a special record. It is created only by a Write Record Zero command and it can be read in its entirety only by a Read Record Zero command. R0 is not preceded by an address mark.

Read Count, Read Count Key and Data, and Write Count Key and Data commands do not operate on R0 (TDR).

Read Key and Data, Read Data, Write Key and Data, Write Data and Search Key operate on R0 (TDR) only if chained from a satisfied Search ID command for R0 (TDR).

The record is divided into three distinct areas:

1. Count area

   - Physical Address

     Two bytes indicating the physical address of the track as in the home address area.

- Flag Byte

  The flag byte for Record Zero (TDR) is copied from the home address flag.

- Cylinder and Head Number

  The four bytes (CCHH) should normally identify the cylinder and head number of the track addressed.

- Record Number

  The one byte record number subfield should contain a unique number (i.e. a record number different than any other on the particular track).

- Key Length

  The key length byte specifies the number of bytes (excluding the 7 ECC bytes) in the key area of the record. If the key length is zero, the key area and its associated gap (G2) are omitted.

- Data Length

  The two data length bytes specify the number of bytes in the data area of the record (excluding the 7 ECC bytes). If the data length bytes are zero:

  a. the data area will contain one byte of zeros plus seven check bytes,

  b. certain commands will post Unit Exception status.

- Error Correction Code (ECC)

  The 7 ECC bytes are generated and written when R0 (TDR) is formatted. The ECC bytes are verified whenever R0 (TDR) is read, searched or clocked (read internally but not transferred to the channel).

2. Key area

   The key area is made up of the user specified key field plus the 7 bytes of ECC. The key field is used to identify the information recorded in the data area of a record. If a key area is rewritten the data area is rewritten as well. The key area is followed by a gap (G2).

   The length of this area is defined by the key length (KL) plus seven (for the ECC bytes). Once the area is formatted, its contents, (but not its length) can be altered. If the key length is zero, the key area and its associated gap (G2) are omitted.

   Seven bytes of ECC are appended to the key field when it is written. The ECC bytes are verified whenever the key is read or searched.

3. Data area

   The data area consists of the user specified data field plus seven ECC bytes. The data field contains user data identified by the count and key fields.

   The length of the data area is defined as the data length (DL) specified in the count field plus seven (for the ECC bytes). An exception occurs when DL is zero, in which case the data field contains one byte of zeros plus seven ECC bytes. If a zero length data field is encountered, certain commands will post Unit Exception status.

   Seven bytes of ECC are written when the data field is written. The ECC bytes are verified whenever the data field is read.

2.4.2.2.4. Data Records (R_1–R_N)

Data records have the following areas:

1. Address Mark

   A three-byte address mark is written by the SCU to identify the beginning of records on the disc unit (see Figure 2–10).

2. Count Area

   The count area is made up of:

PA	F	CCHHR	KL	DL	ECC	
2	1	5	1	2	7	No. of Bytes

   The count field is defined as the 8 byte CCHHRKLDL subarea.

   - The physical address (PA) is as defined for the home address area.

     The flag byte (F) is copied from the home address flag byte with the following exception: Bit 4, if set, indicates that the logical record continues on the next track (i.e. this record has been formatted by the Write Special Count Key and Data command).

   - Cylinder and Head Number

     These CCHH bytes normally indicate the seek address of the track.

   - Record Number

     The record number (R) should be a unique number for each record on a track.

   - Key Length (KL)

     Same as for Record Zero (TDR)

   - Data Length (DL)

     Same as for Record Zero (TDR)

   - Error Correction Code (ECC)

3. Key Area

   Same as for Record Zero (TDR)

4. Data Area

   Same as for Record Zero (TDR)

2.4.2.3. Storage Capacity

Because each record has some nondata characters, such as its count field, the net data storage capacity of the tracks may vary.

The actual number of equal length records on a track may be calculated from:

$$\text{Records/track} = \frac{13,165}{133+C+KL+DL}$$

where:

KL = Key Length (in bytes)

DL = Data Length (in bytes)

C = 0 if KL = 0, or 56 if KL ≠ 0.

NOTE:

This formula is valid for the 8430/8433 discs only.

A major loss of storage area is due to gaps between records. Table 2–4 illustrates the number of equal length records that can be written on a track following home address and standard R0 (TDR).

2.4.3. 8405 Disc Stack

The 8405 nonremovable disc stack contains either six discs per stack (8405–00/01 disc) or three discs per stack (8405–04/05). Each 8405 FHD contains only one disc stack. The disc stack is permanently mounted to the drive shaft. The discs are 14 inches (nominal) in diameter. The track spacing is 0.025 inch and the pulse density is 2,650 pulses per inch (maximum).

Table 2–4. SPERRY UNIVAC 8430/8433 Disc Subsystem Data Byte Capacity for Each Record

Maximum Bytes per Record Formatted without Keys	Equal Length Records per Track	Maximum Bytes per Record Formatted with Keys
13030	1	12974
6447	2	6391
4253	3	4197
3156	4	3100
2498	5	2442
2059	6	2003
1745	7	1689
1510	8	1454
1327	9	1271
1181	10	1125
1061	11	1005
962	12	906
877	13	821
805	14	749
742	15	686
687	16	631
639	17	583
596	18	540
557	19	501
523	20	467
491	21	435
463	22	407
437	23	381
413	24	357
391	25	335
371	26	315
352	27	296
335	28	279
318	29	262
303	30	247
289	31	233
276	32	220
263	33	207
252	34	196
241	35	185
230	36	174
220	37	164
211	28	155
202	39	146
194	40	138
186	41	130
178	42	122
171	43	115
164	44	108
157	45	101
151	46	95
145	47	89
139	48	83
133	49	77
128	50	72

2.4.4. Storage Control Unit (SCU)

The 5039 Storage Control Unit (SCU) controls all the functions for the direct access storage operations. In operation, the SCU accepts selection from one of up to four channels, establishes connection to the requested disc unit, decodes the command from the channel, and controls the transfer of data to or from the disc unit. The operations performed by the SCU are specified by the microprogram instructions which are resident in the nonvolatile Control Read-Only Memory (CROM). The sequential instructions that are read out of CROM are used in place of hard-wired sequential logic.

The SCU controls one to eight 8430/8433 discs. Any 8430 disc can be replaced with an 8433 disc. The 16 Drive Expansion feature allows control of one to eight additional 8430/8433 discs. The 8405 Capability feature allows control of one to eight 8405 discs. These two features are mutually exclusive, i.e., only one of these features may be installed in the SCU at one time.

A microprogram feature is provided for attachment of the SCU to a SPERRY UNIVAC 1100 Series System. This feature is necessary when a buffered SCU is used.

In order to check the accuracy of data transferred to and from the disc unit, the SCU employs extensive error detection and correction facilities. The microprogram utilizes error correction codes (ECC) together with a read/write buffer to temporarily store information about errors. The buffer also contains usage and error counters (Environmental Data) on each disc unit as well as temporary storage of sense data.

In particular, an SCU:

- assumes and relinquishes control of the processor/disc subsystem interface,
- accepts and decodes commands from the processor I/O channel,
- initiates an operation at an addressed disc unit,
- provides a byte parallel to bit serial data conversion,
- ensures correct data transfer by a combination of odd parity checking and error correction codes,
- provides information to the system on the status of all operations,
- provides automatic error recovery techniques to the system,
- provides statistical usage and error data to the system and,
- converts 36-bit words to 8-bit bytes on output from the processor I/O channel and 8-bit bytes to 36-bit words for input to the I/O channel.

3. Programming

3.1. GENERAL

This section provides the programmer with the executable commands, status bytes, sense data bytes, programming considerations and error recovery procedures used in the SPERRY UNIVAC 8405/8430/8433 Disc Subsystem (disc subsystem).

3.2. COMMANDS

The disc subsystem recognizes as valid the commands listed in Table 3–1. Each command is checked for parity and validity before being accepted by the storage control unit (SCU). Status signals generated in response to the commands are normally presented two times:

1. Initial status

   Initial status is presented after the command has been decoded, but before any data or command arguments have been transferred. Initial status contains either zeros, pending, or stacked status for the selected disc unit or SCU. Any nonzero initial status, except for the Test I/O and No Operation commands, indicates that the command was not executed.

2. Ending status

   Ending status is presented for every command, other than Test I/O and No Operation (for which only initial status is given), after:

   - data transfer (if any) is complete and

   - the SCU has carried out all of its requirements for execution of the command.

A complete discussion of status presentation is given in 3.3.

The proper execution of many commands (e.g. Write commands) require precise orientation. Orientation may be viewed as synchronizing the disc unit to the SCU and the channel.

The commands fall into five major categories (see Table 3–1).

Table 3–1. Command Code Summary

	Command	Command Code			
		Multiple Track OFF		Multiple Track ON (If Applicable)	
		Hexa-decimal	Binary	Hexa-decimal	Binary
CONTROL	Seek	07	0000 0111		
	Seek Cylinder	0B	0000 1011		
	Seek Head	1B	0001 1011		
	Set Sector	23	0010 0011		
	Seek and Set Sector	27	0010 0111		
	Recalibrate	13	0001 0011		
	Set File Mask	1F	0001 1111		
	Space Count	0F	0000 1111		
	Retry Restart	3B	0011 1011		
	No Operation	03	0000 0011		
	Restore	17	0001 0111		
WRITE	Home Address	19	0001 1001		
	Record Zero	15	0001 0101		
	Erase	11	0001 0001		
	Count, Key and Data	1D	0001 1101		
	Special Count, Key and Data	01	0000 0001		
	Data	05	0000 0101		
	Key and Data	0D	0000 1101		
SEARCH	Home Address (Equal)	39	0011 1001	B9	1011 1001
	Identifier Equal	31	0011 0001	B1	1011 0001
	Identifier High	51	0101 0001	D1	1101 0001
	Identifier Equal or High	71	0111 0001	F1	1111 0001
	Key Equal	29	0010 1001	A9	1010 1001
	Key High	49	0100 1001	C9	1100 1001
	Key Equal or High	69	0110 1001	E9	1110 1001
READ	Home Address	1A	0001 1010	9A	1001 1010
	Count	12	0001 0010	92	1001 0010
	Record Zero	16	0001 0110	96	1001 0110
	Count, Key and Data	1E	0001 1110	9E	1001 1110
	Key and Data	0E	0000 1110	8E	1000 1110
	Data	06	0000 0110	86	1000 0110
	IPL	02	0000 0010		
	Sector	22	0010 0010		
SENSE	Sense I/O	04	0000 0100		
	Device Release	94	1001 0100		
	Device Reserve	B4	1011 0100		
	Test I/O	00	0000 0000		

3.2.1. Control Commands

Control commands set or reset states in the SCU, perform track addressing and other special functions.

- Control commands which require arguments from the channel generally check the validity of the argument both for content and for sufficient length. If the channel can pass more bytes than required, the excess is not transferred to the SCU.

- Invalid arguments are flagged when Unit Check (Command Reject), Channel End, and Device End are presented in ending status. Control commands (except Space Count and Retry Restart which are self-orienting) reset orientation.

- Error conditions common to all commands

  – Bus Out Parity on the command byte results in Unit Check alone in initial status.

  – Bus Out Parity on the command argument results in Unit Check, Channel End, and Device End in ending status.

  – If the argument is invalid in either content or length, Channel End, Device End, and Unit Check (Command Reject) are posted in ending status.

3.2.1.1. Seek (07_{16})

The Seek or Seek Cylinder command addresses the cylinder and head of the selected disc unit. The SCU requests a six-byte argument from the channel whose format is:

B_1B_2	C_1C_2	H_1H_2

The B_1B_2 bytes of the Seek or Seek Cylinder commands are defined as follows:

- B_1 – PPQXXXX

  PP=11 – The SCU assumed that B_2 is a sector number as defined under the Set Sector command.

  PP=00 – The remainder of B_1 and all of B_2 are ignored.

  PP=$\begin{smallmatrix}01\\10\end{smallmatrix}$ – Channel End, Device End, and Unit Check (Command Reject) are posted in ending status.

  Q=1 – The sector decrement and statistical integrator are enabled and used as in the Seek and Set Sector command.

  Q=0 – The sector decrement and statistical integrator are not enabled and not used. The fixed sector advance is two as in the Set Sector commands.

  X – Ignored

- B_2 – B_2 is meaningful only if PP of B_1=11. Then B_2 is the sector number as defined for the Set Sector commands.

  The operation of the Seek or Seek Cylinder command with B_1 (PPQ=111) is identical to the Seek and Set Sector (X′ 27_{16}) command.

  If B_1 PPQ is set to 110, the sector argument is decreased by two (modulo 128) and issued to the disc unit. After Device End is presented, no count down statistical integrator update is performed.

- C_1C_2 defines the cylinder address. It can be considered to be an unsigned 16-bit binary number where:

8405–04	$0 \leq C_1C_2 \leq 35$
8405–00	$0 \leq C_1C_2 \leq 71$
8430	$0 \leq C_1C_2 \leq 410$
8433	$0 \leq C_1C_2 \leq 814$

- H_1 must be zero.

- H_2 defines the head address, where:

8405–00/04	$0 \leq H_2 \leq 11$
8430	$0 \leq H_2 \leq 18$
8433	$0 \leq H_2 \leq 18$

A seek argument of less than six bytes or not confined to the values above receives Unit Check (Command Reject).

Initial status is normally zero. If access motion is required to address the cylinder, Channel End is presented in ending status. Device End is presented after the access motion is complete and the specified head selected.

There are no chaining requirements for the Seek command.

If a Seek error is detected after presenting Channel End, but before presenting Device End, the SCU will post Device End as normal. Unit Check is presented on the next command to the disc unit.

If a Seek command follows a Set File Mask command and it violates the file mask set, Unit Check (File Protect) is set in initial status.

3.2.1.2. Seek Cylinder (OB_{16})

The Seek Cylinder command functions in exactly the same manner as the Seek command.

3.2.1.3. Seek Head ($1B_{16}$)

The Seek Head command selects the specified head (if valid) without disturbing the cylinder address.

Six bytes are requested from the channel having the format:

$B_1 B_2 C_1 C_2 H_1 H_2$

- $B_1 B_2$ must be zero.
- $C_1 C_2 H_1 H_2$ must be valid in the same ranges as specified for the Seek command.
- The entire argument is checked for validity in both content and minimum length.
- All but H_2 are ignored.

Initial status is normally zero. Ending status is normally Channel End and Device End. Unit Check (Command Reject) is presented with Channel End and Device End if the argument is invalid.

3.2.1.4. Set Sector (23_{16})

The Set Sector command conditions the SCU and disc unit to present a Device End interrupt when the drive reaches a specified angular position.

A one-byte sector argument is requested from the channel. The sector argument may vary between 0 and 127 where 0 corresponds to index and 127 corresponds to the sector before index. The sector argument normally corresponds to the sector before index. The sector argument normally corresponds to the sector which contains the count field of the record to be processed. A sector argument of 255 is allowed and is interpreted as no operation. On the 8430/8433 disc a sector spans 105 bytes and on the 8405 FHD a sector spans 81 bytes.

Initial status is normally zero. Ending status following acceptance of argument byte (n) is:

- Channel End for a sector value $0 \leqslant n \leqslant 127$
- Channel End and Device End for $n = 255$
- Channel End, Device End, and Unit Check (Command Reject) for $127 \leqslant n \leqslant 255$

The SCU decreases the sector argument ($0 \leqslant n \leqslant 127$) by two (modulo 128) before issuing the Set Sector command to the disc unit. The two sector advance provides a reselection and chaining time gap between Device End (when the drive reaches the decremented angular position) and the next command. At reconnection the SCU is "unoriented". Search commands are required to guarantee record positioning and orientation. Multi-track Search commands should be preceded by Set Sector (0) followed by Read Home Address.

3.2.1.4.1. MSA Channel Operation

For an unchained APS Seek, the channel disconnects after Channel End and the SCU attempts to send Device End on the first and subsequent rotations (subject to the one sector request-in time). When Device End is accepted, the I/O sequence is complete.

3.2.1.4.2. Sector Calculation

Record Zero (TDR) begins in byte 155 from index. Sector 0 must be used for Home Address and R0 (TDR) because index orientation is required for both. The sector number (S) for any record R_n ($n \geqslant 1$) from an arbitrary R0 is:

$$S = \frac{155 + \sum_{i=0}^{n-1} \left\{ 133 + \binom{KL_i + 56}{0} + DL_i \right\}}{BS}$$

where:

S	=	sector number
KL	=	key length
DL	=	data length
BS	=	bytes/sector (81 for the 8405 disc, 105 for the 8430/8433 disc)

3.2.1.5. Seek and Set Sector (27_{16})

The Seek and Set Sector command combines the functions of Seek and Set Sector, and includes functions to adjust angular positioning advance in the MSA channel environment. The adjustments which are performed are based on a sector decrement field and a statistical integrator field.

The SCU requests seven bytes from the channel. The format of the bytes is:

$B_1 B_2 C_1 C_2 H_1 H_2 S$

- $B_1 B_2$ must be zero.
- $C_1 C_2 H_1 H_2$ are as defined for the Seek command.
- S is the argument as defined for the Set Sector command.

The entire argument is checked for validity as for Seek and Set Sector commands. An argument of less than seven bytes receives Unit Check (Command Reject). If S is 255, the command is identical in operation to Seek.

3.2.1.5.1. Programming Notes

Initial status is normally zero. Channel End is presented after the seven byte argument is accepted. Channel End and Device End are presented if no access motion is required and the sector argument is 255. Device End is presented when the drive has positioned (addressed) the specified cylinder, selected the specified head and reached the specified sector.

The sector argument for the Seek and Set Sector commands is decremented by five (plus the contents of the sector decrement field, see 3.2.1.5.2). The fixed number five and the variable sector decrement field compensate for system interrupt processing and restart.

Seek and Set Sector should be issued unchained with processor software performing the block multiplexer function of reconnection by starting a new command sequence.

The restart command chain should begin with a Seek command to guarantee cylinder and head addressing for a dual-access system.

Like Set Sector, the Seek and Set Sector command does not guarantee orientation. Search commands must be used for that purpose. In cases where the Search command itself must be oriented (multi-track), the sector argument should be zero and the restart Seek command should be followed by the Read Home Address command.

Any error conditions for the Seek or Set Sector commands are applicable to the Seek and Set Sector command.

3.2.1.5.2. Sector Decrement and Statistical Integrator Fields

The addition of 64 or subtraction of one is performed on the entire 16-bit field (shown below) such that the sector decrement field (activated by Q=1, see 3.2.1.1) is updated by overflow from or underflow to the statistical integrator. A power-on reset of the SCU clears the 16-bit field. There is a sector decrement and statistical integrator associated with each channel interface.

0 — 7	8 — 15
Sector Decrement	Statistical Integrator

When the channel accepts Device End from an unchained Seek and Set Sector, the SCU begins a timing count down of one sector time multiplied by one plus the variable number in the sector decrement field. During the count down the SCU is busy to all channel interfaces except the one which issued the Seek and Set Sector command.

Where the next command is not addressed to the appropriate drive, Control Unit Busy followed by Control Unit End is presented. The statistical integrator is reduced by one and the command is executed.

If a command is issued from the appropriate channel to the appropriate drive before the count down has completed, the "reconnection" is assumed to have occurred properly. The statistical integrator field is decreased by one and the new command chain proceeds. If 256 successful "reconnections" have occurred the sector decrement field is reduced by one such that the variable advance for the next Seek and Set Sector command is one less than the previous one.

If the count down terminates before a command is issued from the appropraite channel and to the appropriate drive, the statistical integrator field is increased by 64. If four unsuccessful "reconnections" have occurred, the sector decrement field is increased by one such that the variable advance for the next Seek and Set Sector is one greater than the previous one. After incrementation, the SCU is no longer busy to any channel interface.

The values in the sector decrement and statistical integrator fields are displayed in sense bytes 22 and 23 respectively for formats 0, 1, and 2 sense.

3.2.1.6. Recalibrate (13_{16})

The Recalibrate command causes the selected disc unit to seek to cylinder zero and head zero. On the 8430/8433 the access arm is first moved to the inner guard band (beyond the innermost recording cylinder) before the seek to zero. No argument is requested from the channel.

Initial status is normally zero. Ending status is:

- Channel End and Device End on the 8405 FHD and
- Channel End only on the 8430/8433 disc.

Device End is presented after access motion is complete on the 8430/8433 disc.

Recalibrate is subject to the same error conditions as Seek including file mask violations.

3.2.1.7. Set File Mask ($1F_{16}$)

The Set File Mask command defines limits of types of seeks, writes, and can also be used to enable or inhibit certain SCU error recovery options. Once set, the file mask remains in effect throughout the remainder of the command chain. The file mask has a default value of zero at the beginning of every command chain.

One byte of data (the file mask) is requrested from the channel. The file mask byte format is:

Bit	0	1	2	3	4	5	6	7
	W	W	O	S	S	O	T	P

- WW – 00 – Inhibit Write Home Address and Write Record Zero (TDR)
  01 – Inhibit all Write commands
  10 – Inhibit all Format Write commands
  11 – Allow all Write commands
- Bit 2 must be zero.
- SS – 00 – Permit all Seek commands
  01 – Permit Seek Cylinder and Seek Head commands
  10 – Permit Seek Head commands only (multi-track and overflow record operations are allowed)
  11 – Inhibit all Seeks including multi-track and overflow record head switches.
- Bit 5 must be zero.
- T – 0 – Inhibit transparent Command Retry
  1 – Allow transparent Command Retry
- P – 0 – Normal Command Retry operation

  1 – If Command Retry has been used to correct an error, the SCU will generate Unit Check (Data Check) in initial status for the retried command. The Data Check may be a dummy (all zero correction bytes). The intent of the Unit Check is only to interrupt the processor.

Initial status is normally zero. Ending status is normally Channel End and Device End. Unit Check (Command Reject) is presented with Channel and Device End if the argument is invalid.

Some systems may use command chains in such a manner that data is read and processed before the chain terminates. Such systems are identified by the P-bit and the SCU interrupts the processor at the point of a corrected (retried) error. Only one Set File Mask may appear in a chian. Unit Check (Command Reject) is posted in initial status if a second one is encountered.

3.2.1.8. Space Count ($0F_{16}$)

The Space Count command is a special data recovery command which may be used to retrieve the key and data fields following an uncorrectable count field. The defective count field is not read or verified and the required control information usually obtained from the count field must be supplied as the argument to the Space Count command.

Three bytes are requested from the channel. The format of the argument is:

KL (Key length)	DL (Data length)	DL (Data length)
Byte 1	Byte 2	Byte 3

No validity checking is done for contents nor minimum length of the argument. Missing bytes are assumed to be zero.

If the Space Count is received while the SCU is unoriented (e.g. after a Control command) the SCU orients to the Record Zero (TDR) count area, skips it and presents Channel End and Device End.

If the Space Count is received while the SCU is oriented (e.g. after a Search, Read or Write command), the SCU orients to the next count area, skips it and presents Channel End and Device End.

The internal file mask is modified to inhibit all Write commands. Write or Set File Mask commands are rejected if chained from a Space Count command.

3.2.1.9. Retry Restart ($3B_{16}$)

The Retry Restart command is a special command which can be used to recover from a data check. The Retry Restart command has as its argument the entire 24 bytes of Format 7 Sense information received from a previous sense I/O.

The Format 7 Sense contains all the information required to restore the SCU to the content of the error and prepare it to retry the failing command. In addition, the Format 7 message code describes the specific nature of the error.

Operationally the Retry Restart command performs a Seek, a Set Sector, restores the file mask, and orients to the proper bytes for the retry. Up to 27 sequences of Format 7 Sense and retry restarts can occur, during which the SCU is performing recovery operations. If the error is unrecoverable, the SCU posts Unit Check (Data Check and Permanent Error) after 27 attempts. If the error is corrected internally, (error in home address or count fields) the command chain will continue normally. If the data check becomes correctable (key or data field error) the SCU posts Unit Check (Data Check; Correctable) with sufficient information in the sense to correct the error.

The Retry Restart command must be the first command of a chain and it must be chained to the command which encountered the Unit Check which produced the Format 7 Sense. Unit Check (Command Reject) is posted if the Retry Restart command is not the first command in a chain. Initial status is normally zero.

Channel End is presented when the 24-byte transfer is complete. Unit Check (Command Reject) with Channel End and Device End is presented if:

- The Format 7 Sense sent back to the SCU is invalid.
- Less than 24 bytes of sense is received by the SCU.
- The physical drive identification specified in the Format 7 Sense is not the selected drive.
- The Retry Restart command is not chained.

Device End is presented when orientation is complete.

The command following the Retry Restart command receives Unit Check (Command Reject) if the command code is unequal to the failing command code.

Between the Sense I/O command which received the Format 7 Sense, and the Retry Restart command; the disc subsystem may be used for other operations. If the disc unit having the error has the dual-access feature, the retry can be performed through the alternate SCU if both accesses have the same physical address (3 of 6 code).

If the module select plug (logical address) has been changed between the sense and the restart, the restart should be performed by addressing the current (new) logical address. Presenting an argument which is not totally a Format 7 Sense, but not detectably invalid may render the SCU inoperable until it is manually reset.

3.2.1.10. No Operation (03_{16})

The No Operation command causes no operation. Initial status is normally Channel End and Device End.

No argument is requested.

3.2.1.11. Restore (17_{16})

The Restore command operates quite similar to the No Operation command. The major difference is that Channel End and Device End are presented in ending status. Initial status is normally zero.

No argument is associated with the Restore command.

3.2.2. Write Commands

Write commands are used to record information on the storage medium. The Write command set is divided into two types:

- Format writes
- Update writes

The prerequisites for performing any Write commands are very stringent and any violation results in Unit Check (Command Reject) in initial status. The prerequisites are:

1. Bits 0 and 1 of the file mask are defined to allow or disallow the execution of Write commands.

Bit 0	Bit 1	Description
0	0	Inhibit write HA and R0 (TDR) (allow all other writes)
0	1	Inhibit all Write commands
1	0	Inhibit all Format Write commands
1	1	Permit all Write commands

2. The READ ONLY (write inhibit) switch on the 8430/8433 discs must not be set to the Read Only position.

3. No Write command can be the first command in a chain. The particular chaining requirements vary and will be presented later.

4. If the Write command is not preceded by a Set File Mask command, the default File Mask (00_{16}) command is in effect.

Bytes accepted from the channel are not checked for validity. Data bytes in excess of the number required are not transferred to the SCU. If the channel truncates, the remainder of any fields is padded with zeros.

Initial status is normally zero. Ending status is normally Channel End and Device End, presented after the ECC bytes for the last field defined by the command have been written. An error-free Write command does not guarantee that the information recorded is readable.

3.2.2.1. Error Conditions

The error conditions common to all Write commands are:

1. Bus Out (channel transmits output data) Parity error on the command bytes results in Unit Check in initial status.

2. Bus Out Parity error on a data byte results in Unit Check with Channel End and Device End in ending status. Ending status is presented after the ECC bytes have been written for the current field. The remainder of the field from the point of error is padded with zeros. Format writes erase to index after ending status is presented.

3. Command Reject (Unit Check) is presented in initial status if:

   - The write portion of the file mask is violated
   - The READ ONLY (write inhibit) switch is enabled
   - Chaining requirements are not met

4. Command Overrun results in Unit Check in initial status

3.2.2.2. Types of Write Commands

A discussion of the two types of Write commands follows.

3.2.2.2.1. Format Writes

The following commands are classed as format writes:

1. Write Home Address (19_{16})
2. Write Record Zero (15_{16})
3. Erase (11_{16})
4. Write Count, Key, Data ($1D_{16}$)
5. Write Special Count, Key, and Data (01_{16})

Format writes record control information (home address and count fields) in addition to recording key and data areas. The count field, as it is presented by the channel, determines the record key and data field lengths. Format writes cause the remainder of the track following the last field written to be erased.

In order to maintain track format validity, the SCU erases the remainder of the track after the ECC bytes have been written for the last field defined by the command. However, format writes can be chained together without intervening searches and the erase does not occur until the last field defined by the last format write has been written. Following the erase the SCU is not oriented.

If the command following a format write is not a format write, the execution is delayed until the remainder of the track has been erased.

If the chain terminates on a format write and another chain is started while the erase is being performed, the SCU responds by presenting Control Unit Busy (Status Modifier and Busy). Control Unit End will be presented at index.

Data overrun results in Channel End, Device End and Unit Check in ending status after the ECC bytes have been written for the last field defined by the command. The record is padded with zeros from the point of the detection of the error to the end of the last field defined. The remainder of the track is erased.

3.2.2.2.2. Update Writes

The following commands are classed as update writes:

1. Write Data (05_{16})
2. Write Key and Data ($0D_{16}$)

Update writes modify the information contained in the key or data areas, but do not change their lengths. The SCU reads the count field of the record and uses the record descriptors (key length and data length) to control the operation. An update write performed on an overflow record can modify fields in more than one physical record.

Update writes are used to replace data or key and data fields previously written. The count field (CCHHRKLDL) of the record to be updated is read and the KLDL subfields are used to control the writing.

Update writes must be chained from a satisfied Search Identifier Equal or Search Key Equal command where the channel did not truncate the search argument.

NOTE:

The truncated (including zero length) search arguments (Search ID Equal only) allow writes if a previous full Search ID Equal was satisfied on a non-truncated argument and the SCU has remained in the oriented state since that time.

In addition to the general Write error conditions, update writes have the following error conditions:

1. Data overrun

   Data overrun results in Channel End, Device End, and Unit Check (overrun) being presented in ending status. The remainder of the record or segment is padded with zeros.

2. Command overrun

   Command overrun occurs if the next command is late with respect to the next field on the disc unit. A Write Key and Data command following a Search Identifier Equal command must be received before any of the key and data fields of the record have been passed.

   Channel End, Device End, and Unit Check (overrun) are presented in ending status.

3. Zero length record

   Channel End, Device End and Unit Exception are presented in ending status if an attempt is made to update a record having a zero data length.

3.2.2.3. Write Home Address (19_{16})

The SCU orients to index and requests five bytes from the channel interface to define FCCHH. The physical address is written and FCCHH and its associated ECC bytes are written. If the channel truncates the requested field bytes, the remainder of the field is padded with zeros. If the command following the Write Home Address command is not Write Record Zero the remainder of the track is erased.

The Write Home Address command must be preceded by a Set File Mask command with the file mask bits (0,1) set to binary (11).

3.2.2.4. Write Record Zero (Write Track Descriptor Record) (15_{16})

The Write Record Zero command is written as the first record after the Home Address command. The first eight bytes requested from the channel define the count field (CCHHRKLDL) of Record Zero (TDR). The flag byte is copied by the SCU from the home address flag. KL and DL define the key and data areas to be formatted. Eight plus KL plus DL bytes are requested from the channel. If the channel truncates, the remainder of any field is padded with zeros.

The Write Record Zero command must be chained from a Write Home Address or Search Home Address command and be preceded by a Set File Mask command with file mask bits (0,1) set to binary (11).

3.2.2.5. Erase (11_{16})

The SCU orients as if to perform a Write Count, Key and Data, but then erases the remainder of the track. The first eight bytes from the channel are interpreted as a count field, but no record is written. Key and data bytes are requested from the channel, but none are written. (The number of key and data bytes are determined by the pseudo count field as in Write Count, Key and Data.)

The command prerequisites are the same as for the Write Count, Key and Data command. The command is useful in determining remaining track capacity.

Unit Check (Command Reject) initial status is given on any format write command which is chained to an Erase command.

3.2.2.6. Write Count, Key and Data ($1D_{16}$)

The first eight bytes requested from the channel define the count field (CCHHRKLDL) of the record written. The flag byte is copied from the home address or a previous count field. The count field in turn defines the key and data areas to be formatted. After writing the count area (including the ECC bytes) the SCU requests KL plus DL additional bytes from the channel, which if available, are written into the key and data fields. If the channel truncates, the remainder of any field is padded with zeros.

The Write Count, Key and Data command must be chained from:

1. A Write Record Zero command

2. A Write Count, Key and Data command

3. A satisfied Search Identifier Equal command where the search argument was not truncated by the channel

   NOTE:

   Writes after truncated searches are allowed if a prior five byte Search Identifier Equal command was satisfied and the SCU has remained oriented since the search was satisfied.

4. A satisfied Search Key Equal command where the search argument was not truncated by the channel.

The default file mask (no Set File Mask command issued) is sufficient to allow the Write Count, Key and Data command.

3.2.2.7. Write Special Count, Key and Data (01_{16})

The Write Special Count, Key and Data command is identical in operation to the Write Count Key, and Data command, except flag byte bit 4 is set to 1 indicating an overflow record.

3.2.2.8. Write Data (05_{16})

The data area of the record is rewritten. The data field is filled with bytes requested from the channel. The number of bytes requested is contained in the DL subfield of the count of the updated record. If the channel truncates, the remainder of the data field is padded with zeros. The data field of Record Zero can be updated under the default file mask if a Search Identifier Equal command for Record Zero has been satisfied.

3.2.2.9. Write Key and Data ($0D_{16}$)

The key and data areas of the record are rewritten. The key and data fields are filled by bytes requested from the channel. The number of bytes requested is (KL+DL). If no key field is presented, the command functions as a Write Data command. The Write Key and Data command must be chained from a satisfied Search Identifier Equal command only. The key and data fields of Record Zero can be updated under the default file mask if a Search Identifier Equal command for Record Zero has been satisfied.

3.2.3. Search Commands

Information from the channel (the search argument) is logically compared to information from the selected disc unit and head. The command code specifies both the field to be searched and the type of comparison to be made.

When orientation to the next field of the type specified is reached, bytes read from the disc unit are logically compared (as specified by the command) to those read from the channel. Ending status of Channel End and Device End is presented after verification of the ECC bytes following the field searched. If the logical comparison specified by the command was satisfied, Status Modifier is presented in ending status together with Channel End and Device End.

No validity checking, other than for parity, is performed on a search argument. Bytes in excess over the number in the field being searched are not transferred to the SCU.

The logical comparison is assumed to be satisfied for the bytes not compared. Ending status will include Status Modifier for a zero length search argument. For truncated (nonzero length) search arguments, the SCU will present Status Modifier if the logical comparison is satisfied for the number of bytes available.

There are no chaining requirements for Search commands.

Initial status is normally zero.

If the multi-track bit (bit 0) of the command code is set, the operation is identical to a non-multi-track search except in the case where index is detected during the orientation to the next field. At the detection of index, the SCU will increase the head address by one and the orientation will continue on the next track. The oriented state is maintained over a multi-track head switch. To ensure that the first track is completely searched before the head switch, a Read Home Address command should precede a multi-track search.

The error conditions common to all Search commands are:

1. A Bus Out Parity error on the command bytes results in Unit Check alone in initial status

2. A Bus Out Parity error on the search argument results in Unit Check with Channel End and Device End in ending status

3. A Data Check results in a Command Retry situation

4. If, on a non-multi-track search, index is detected for the second time within a string of Search commands (without an intervening Read Home Address command, read of a data field or a Control, Sense, or Write commands) a No Record Found condition results. Unit Check, Channel End, and Device End are presented in ending status.

5. If index is encountered during a multi-track search and the file mask prohibits all seeks, the SCU does not increase the head address and a File Protect error results in Unit Check in ending status with Channel End and Device End.

6. If index is encountered on the last track of a cylinder during a multi-track search, the head address is not increased and an End of Cylinder condition results in Unit Check in ending status with Channel End and Device End.

7. If, during a Search ID or Key, the SCU detects a defective track bit in the flag byte of a count field, track condition posts Unit Check with Channel End and Device End in ending status.

8. A command overrun or data overrun error results in Channel End, Device End and Unit Check in ending status.

3.2.3.1. Search Home Address (39_{16} or $B9_{16}$)

When orientation to index is reached, the search argument is logically compared for equality to the four byte CCHH home address field. If the compared bytes are equal, Channel End, Device End, and Status Modifier are presented in ending status. Otherwise, Channel End and Device End appear alone. The Search Home Address command ignores the defective track bit in the flag byte. A multi-track Search Home Address command always attempts a head switch during orientation, hence if the head switch is allowed, the comparison is performed on the next track.

3.2.3.2. Search Identifier (ID) Equal (31_{16} or $B1_{16}$)

After orientation to the next count area is reached (including Record Zero (TDR)) the search argument is logically compared for equality to the five byte identifier field (CCHHR). If the compared bytes are equal, Status Modifier is presented with Channel End and Device End in ending status. Otherwise, Channel End and Device End appear alone.

3.2.3.3. Search Identifier (ID) High (51_{16} or $D1_{16}$)

The Search Identifier High command is identical to the Search Identifier Equal command except that the logical comparison is for the ID bytes from the disc unit to be of greater magnitude than the search argument.

3.2.3.4. Search Identifier (ID) Equal or High (71_{16} or $F1_{16}$)

The Search ID Equal or High command is identical to the Search ID Equal command except the logical comparison tested is for the ID bytes from the disc unit to be equal to or greater than the search argument.

3.2.3.5. Search Key Equal (29_{16} or $A9_{16}$)

The Search Key Equal command tests if the argument is logically equal to the recorded key field. The Status Modifier is included with Channel End and Device End for equal comparison.

Unless the Search Key Equal command is chained from a Space Count, Read Count or Search Identifier command, the SCU orients to the next count field before processing the search. A Record Zero key can be searched if chained from a Search Identifier command for Record Zero (TDR).

If the Search Key Equal command is executed on a record having no key, Channel End and Device End are presented in ending status.

3.2.3.6. Search Key High (49_{16} or $C9_{16}$)

The Search Key High command is identical to the Search Key Equal command except the relationship tested is for the recorded key to be logically greater than the search argument.

3.2.3.7. Search Key Equal or High (69_{16} or $E9_{16}$)

The Search Key Equal or High command is identical to the Search Key Equal command, except the relationship tested is for the recorded key to be logically equal to or greater than the search argument.

3.2.4. Read Commands

Read commands are used to retrieve information from the disc unit and transfer it to the channel interface.

Except for Read Sector, all Read commands transmit information recorded on the storage media. Each field transferred is verified by processing the ECC bytes appended to each field. A parity bit is added to each byte going to the channel interface. After a Read (excluding Read Sector), the SCU is oriented.

The data transfer is initiated by the SCU. If the channel truncates, the control unit stops data transfer to the channel, but processes the command normally at the drive interface. Ending status is presented after the ECC bytes have been verified. If the channel is able to accept more bytes than the field or record possesses, the control unit stops transmission when the field or record is exhausted. Ending status is presented after ECC verification.

Read commands other than Read Sector can operate on overflow records.

Except for Read IPL, there are no chaining requirements for Read commands.

Initial status is normally zero. Ending status is normally Channel End and Device End presented together.

Read commands, other than Read Sector and Read IPL, can operate in multi-track mode by setting bit 0 of the command code. The command function remains identical except for the case where index is encountered while the SCU is orienting to the next record or home address. If this occurs, the SCU will increase the head address by one and the read will take place on the next track. The SCU remains oriented after the head switch.

The error conditions common to all Read commands are:

1. Bus Out Parity on the command byte results in Unit Check alone in initial status.

2. No Record Found is caused by the detection of index for the second time without an intervening Read Home Address command, read or write of a data field, Sense or Control command. Channel End, Device End and Unit Check are posted in ending status.

3. File protect exceptions can occur on multi-track commands when the File Mask is set to inhibit all seeks and a head switch is attempted. The head number is not increased and Unit Check, Channel End, and Device End are presented in ending status.

3. File protect exceptions can occur on multi-track commands when the file mask is set to inhibit all seeks and a head switch is attempted. The head number is not increased and Unit Check, Channel End, and Device End are presented in ending status.

4. End of Cylinder exceptions can occur on multi-track reads if the head number is already the maximum when head switching is attempted. Neither the head nor the cylinder is incremented and Unit Check with Channel End and Device End is presented as ending status.

   NOTE:

   On the 8405 FHD only, the End of Cylinder occurs only at the end of the alternate area. At logical End of Cylinder, the head address is reset and the cylinder number is incremented by one.

5. Unit Exception, Channel End and Device End status are posted after key field processing (if any) on Read IPL, Read Record Zero, Read Count, Key and Data, Read Key and Data and Read Data commands if the data length specified in the count field is zero.

6. Operation Incomplete is set if a read of an overflow record is terminated prematurely. Operation Incomplete can be posted alone or as a modifier bit in conjunction with:

   - Data Check
   - File Protect
   - End of Cylinder
   - Track Condition Check

   Unit Check, Channel End, and Device End are posted in ending status.

7. Command Overrun and Data Overrun result in Unit Check, Channel End, and Device End in ending status.

8. Data Check is processed by transparent Command Retry or results in Unit Check, Channel End, and Device End in ending status.

9. Defective Track results in Unit Check, Channel End, and Device End presented in ending status.

3.2.4.1. Read Home Address ($1A_{16}$ or $9A_{16}$)

The five byte home address (FCCHH) field is transmitted from the disc unit to the channel interface.

The SCU orients to index on the selected drive and head and then transmits the field.

In the case of a multi-track Read Home Address command, a track switch always occurs during orientation and the home address field read is that from the following track.

There are no chaining requirements for the Read Home Address command.

3.2.4.2. Read Count (12_{16} or 92_{16})

The eight byte count field (CCHHRKLDL) of the next count area (excluding Record Zero) is transferred from the disc unit to the channel interface.

If in processing Read Count multi-track, index is encountered while orienting to the next count field, the head switch is performed. The count field read is from the first data record of the track.

3.2.4.3. Read Record Zero (Read Track Descriptor Record) (16_{16} or 96_{16})

The SCU orients to index, clocks G1, HA and G2 and transfers the count, key (if any) and data (if any) fields. The number of bytes transferred is 8+KL+DL where KL and DL are from the Record Zero (TDR) count field. If Read Record Zero command is chained from a Home Address command the SCU clocks G2 and begins the data transfer.

In multi-track mode, the Read Record Zero command transmits the Record Zero (TDR) from the next track unless chained from a Read Home Address command.

If DL is zero, UnitException if presented with Channel End and Device End in ending status.

3.2.4.4. Read Count, Key and Data ($1E_{16}$ or $9E_{16}$)

The count, key and data fields of the next record are transferred from the disc unit to the channel interface, 8+KL+DL bytes are transferred. The KL and DL parameters are read from the subject count field. If DL is zero, Unit Exception is included in ending status with Channel End and Device End.

For a multi-track Read Count, Key and Data command the record read is from the next track if index is encountered during orientation to the next record.

3.2.4.5. Read Key and Data ($0E_{16}$ or $8E_{16}$)

The execution of the Read Key and Data command varies depending on whether the SCU is oriented.

1. Unoriented:

   The SCU orients to the next count area, reads the count field (internally), uses the KL and DL parameters and transfers the record key and data fields to the channel.

2. Oriented:

   a. If the SCU is oriented from a Read Count, Space Count or Search Identifier command the key and data fields of the current record are transferred.

   b. If the SCU is oriented from a key or data field command, the SCU orients to the next count field as in the unoriented case.

A Read Key and Data command on a record not having a key functions as a Read Data command.

For a Read Key and Data command with multi-track set, a track switch is done if index is encountered while orienting to the next count field.

Unit Exception, Channel End, and Device End are posted in ending status if DL is zero.

3.2.4.6. Read Data (06_{16} or 86_{16})

The Read Data command causes the transfer of a data field from the disc unit to the channel. The execution of the Read Data command varies depending on whether the SCU is oriented.

1. Unoriented:

   The SCU orients to the next count field and transfers the data field associated with it.

2. Oriented:

   a. If the SCU is oriented from a Read Count, Space Count, Search Identifier or Search Key command, the data field of the current record is transferred.

   b. If the SCU is oriented from a data field operation, the SCU orients to the next count field and transfers the associated data field.

If the DL parameter is zero, Unit Exception is presented in ending status with Channel End and Device End.

A multi-track head switch occurs on a multi-track Read Data command if index is encountered while orienting to the next count field.

3.2.4.7. Read IPL (02_{16})

The Read IPL command causes the SCU to seek to Cylinder 0, Head 0 and read the data field of the first data record (excluding Record Zero). The data transfer is identical to the Read Data command.

Command Reject (Unit Check) initial status is presented if the Read IPL command is preceded by the Set File Mask command.

3.2.4.8. Read Sector (22_{16})

The Read Sector command transmits one byte giving the angular position (relative to index) of the last count area processed by a Read, Write, Space Count or Search command. Home Address and Record Zero (TDR) are referenced by zero.

The range of angular position is 0 to 127.

Execution of a Read Sector command causes a loss of orientation.

3.2.5. Sense Commands

With the exception of Test I/O, all Sense commands attempt to transfer 24 bytes of information. If the channel truncates the transfer, the SCU stops transferring bytes but the command executes normally. All Sense commands cause loss of orientation. Sense data bytes are given in 3.4.

A Sense command is usually issued in response to a Unit Check condition. The sense information obtained details the cause of the Unit Check.

After Unit Check has been accepted by the channel, the SCU enters a contingent connection (contingent connection is a special state that the SCU maintains after presenting Unit Check (without Status Modifier) to the channel to which the error is associated).

The SCU is busy to all channel interfaces except the one which accepted Unit Check. The SCU is busy on that channel for all devices other than the one causing Unit Check. Any command other than No Operation or Test I/O addressed from the appropriate channel to the appropriate disc unit breaks the contingent connection. The Sense I/O command is normally used for this purpose.

Sense information (if any) is reset when the contingent connection is broken.

3.2.5.1. Sense I/O (04_{16})

Twenty-four bytes of sense information are transferred from the SCU to the channel interface. If the Sense command is the first command (other than No Operation or Test I/O) issued to a device which has generated a Unit Check (established a contingent connection), the 24 bytes of sense describe the error. Otherwise, 23 bytes of zeros plus one byte (byte 4) containing the physical device address are transferred.

After a Sense command has been issued and the bytes have been presented to the channel, the sense information (if any) is reset and the contingent connection (if any) is broken.

Initial status is normally zero. Except in the case of a Bus Out error on the command, the Sense I/O command never generates Unit Check; however, a Sense command issued after Unit Check has been presented may be rejected by Control Unit Busy status. This occurs most commonly with errors encountered during format write commands when the SCU is busy performing the erase. The busy condition is temporary.

Ending status is Channel End and Device End.

There are no chaining requirements for the Sense I/O command.

3.2.5.2. Device Release (94_{16})

The Device Release command performs all the functions of the Sense I/O command and in addition causes the release of the selected disc unit. After the release, the SCU presents Device End status to each channel which received Device Busy during the reserved period.

In the case where the disc unit is a dual-access device, the disc unit is released from the reserving SCU and the reserving channel interface. After the release, the disc unit presents a release interrupt to the nonreserving SCU, if that SCU had attempted selection during the reserved period. Both SCUs will present Device End status to channels as outlined above.

If the Device Release command is chained to the Set File Mask command, Unit Check (Command Reject) is posted in initial status.

3.2.5.3. Device Reserve ($B4_{16}$)

In addition to performing the Sense I/O function, the Device Reserve command reserves the device to the channel interface which issued the command.

In the case where the disc unit is a dual-access device, it is reserved to the port of the issuing SCU and to the channel interface of the issuing channel.

After the reserve, any attempt at selection of the reserved device by another channel will receive Device Busy initial status. In the dual-access case, Device Busy is also presented if selection is through the nonreserving SCU.

The Device Reserve command is accepted and executed by any SCU regardless of its features or configurations and it generates Unit Check (Command Reject) if it is preceded by a Set File Mask command in the same chain.

3.2.5.4. Test I/O (00_{16})

The Test I/O command is an "immediate" command. The command is complete after the SCU presents initial status. Its use is to test and present the status of the SCU and/or the selected disc unit.

Initial status (the only status presented for Test I/O) is normally zero. Pending or stacked status for the SCU or disc unit is presented. The Test I/O command does not break a contingent connection nor does it cause the reset of sense information.

3.3. STATUS BYTE

Information concerning acceptance of a command, status of the disc subsystem or individual disc unit, and the performance of the command is presented to the channel in a status byte. After the status byte is presented to the channel and is accepted, the status byte is cleared.

The bits of the status byte are illustrated in Figure 3–1.

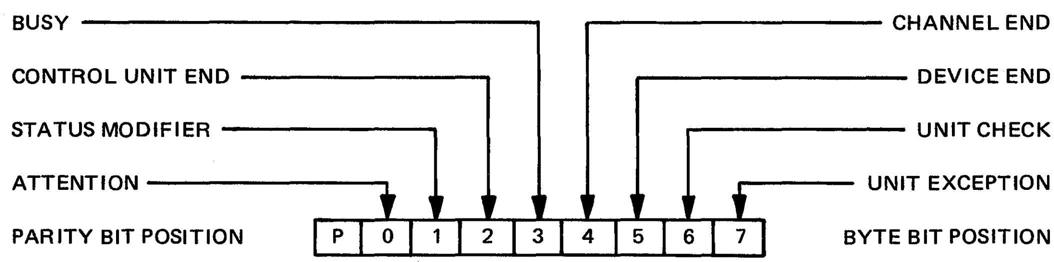

Figure 3–1. Status Byte Format

3.3.1. Attention

Attention status indicates the device has gone from a not-ready-to-ready status. Attention is presented when a module select plug is inserted on an otherwise ready device or when the power-on sequence completes with a module select plug already inserted. Attention is presented to all attached channel interfaces. If the device is a dual-access drive, Attention is generated for both SCUs and all associated channels. If Attention has been stacked it is presented alone in initial status for the Test I/O command or presented with Busy status for any command other than Test I/O.

NOTE:

Attention can be presented as an unsolicited status, indicating the disc unit is ready for operation.

3.3.2. Status Modifier

Status Modifier status is presented when:

- Status Modifier with Busy indicates the SCU is busy
- Status Modifier with Channel End and Device End indicates that a Search command comparison is satisfied

3.3.3. Control Unit End

Control Unit End status is a non-command associated response to Control Unit Busy (Status Modifier with Busy). If during initial selection, Control Unit Busy was presented, the SCU when free, presents Control Unit End as unsolicited status to the channels which received Control Unit Busy.

The device address associated with Control Unit End can be any address recognizable to the SCU, regardless of whether the device is attached or ready.

3.3.4. Busy

Busy status is presented only in initial status. Busy alone implies that the selected device is busy. Busy with Status Modifier indicates Control Unit Busy. When the disc unit or SCU becomes free, Device End or Control Unit End respectively are presented. Device Busy is presented if:

1. The disc unit is performing an operation for which Channel End (but not Device End) has been given
2. The disc unit is reserved to another channel interface

3. A dual-access disc unit is busy on the other SCU

4. A dual-access disc unit is reserved to the other SCU

5. A pending Device End for the addressed disc unit exists, in which case Busy and Device End are presented together

6. A pending or stacked Control Unit End exists, in which case Busy and Control Unit End are presented together

7. Pending or stacked status for the addressed device exists, in which case Busy and the other status are presented together

NOTE:

Cases 5, 6, and 7 above, apply to all commands except Test I/O.

3.3.5. Channel End

Channel End status is normally presented for each command and indicates that data transfer for the command is complete. Commands which have no associated data transfer post Channel End after the operation is started. Channel End implies that Device End will also be presented. The Device End is either included with Channel End or is presented alone at a later time.

Channel End alone is presented for:

1. Seek, Seek Cylinder, Seek and Set Sector, and Recalibrate commands where access arm movement is required

2. Set Sector or Seek and Set Sector commands where rotational delay occurs

3. Retry Restart command

4. Search commands on Transparent Command Retry

5. Seek or Seek Cylinder commands with angular position sensing enabled where rotational delay occurs

3.3.6. Device End

Device End status is presented at the completion of each command at the device. It may appear with Channel End or be presented alone later to the SCU (dual-access) and/or channel which received Channel End.

Device End is also presented in response to Device Busy status (Busy Alone). In this case, Device End is presented when the device becomes free (nonbusy or released).

NOTE:

Device End presentation can be unsolicited in some cases.

3.3.7. Unit Check

Unit Check status is a command associated status and indicates an unusual (error) condition has occurred on either the current or the prior command. When not presented with Status Modifier, the SCU constructs sense information detailing the error condition and enters a contingent connection. A Sense I/O command should be issued to transfer the sense information and break the contingent connection.

Unit Check (without Status Modifier) may appear alone (initial status) or with Channel End and Device End (ending status). Generally, if errors are detected after Channel End has been presented, Device End will be presented as normal and the next command receives the Unit Check in initial status. Unit Check can also occur after Channel End and appear with or without Device End.

NOTE:

Unit Check presentation can be unsolicited in some cases.

3.3.8. Unit Exception

Unit Exception status indicates that an attempt to process a "zero-length" data field has occurred. Unit Exception is generated only on the following commands:

1. Read Record Zero (TDR)
2. Read Count Key and Data
3. Read Key and Data
4. Read Data
5. Read IPL
6. Write Key and Data
7. Write Data

3.3.9. Status Presentation

The status byte is transmitted to the channel in any of the following situations.

1. During the initial-selection sequence
2. To present a Channel End status at the termination of data transfer
3. To present the Device End signal and any associated conditions to the channel. The I/O device remains busy during an operation until the channel accepts the Device End status.
4. To present Control Unit End or Device End status which signals that the SCU or disc unit that was previously selected while busy, is now free. However, the device remains busy until the channel accepts Control Unit End or Device End.
5. To present any previously stacked status when allowed to do so

6. To present any externally initiated status to the channel ('Attention') because of not-ready-to-ready transition

Once accepted by the channel, any given status byte is reset and is not presented again.

3.4. SENSE DATA BYTES

Sense data bytes are developed by the disc subsystem to give detailed information about an error or unusual condition.

The sense information is constructed by the SCU whenever Unit Check status is presented to the channel (without the Status Modifier bit raised). The SCU then enters a contingent connection state to insure that the sense information is preserved.

A sense byte record is 24 bytes long. There are nine different formats for the sense byte records:

1. Format 0 – Programming and some SCU errors
2. Format 1 – Disc unit (device) checks
3. Format 2 – SCU checks
4. Format 3 – Selective reset
5. Format 4 – ECC uncorrectable data errors
6. Format 5 – ECC correctable data errors
7. Format 6 – Error and usage counts
8. Format 7 – Disconnected Command Retry
9. Format F – Inline sense

To determine the specific format, the sense byte record can be broken into three sections:

SECTION 1	SECTION 2	SECTION 3
BYTES 0 6	BYTE 7	BYTES 8 23

- Section 1, bytes 0–6 provide general information about the error or check condition. The interpretation of bytes 0–6 is the same for all formats.

  NOTE:

  Bytes 0–6 have a different interpretation for Format 7 Sense bytes. A detailed discussion of bytes 0–6 starts at 3.4.1.

- Section 2, byte 7 indicates the format (and message number) for sense bytes 0–23. Byte 7 is broken down as follows:

BIT	0 1 2 3	4 5 6 7
	THIS HEXADECIMAL CHARACTER INDICATES THE FORMAT NUMBER	THIS HEXADECIMAL CHARACTER INDICATES THE MESSAGE NUMBER

- Section 3, bytes 8–23 provide detailed information on the type of error or check. The interpretation of these bytes varies according to format.

Sense data bytes and sense formats are discussed in detail in the following paragraphs. Figure 3–2 summarizes the individual bit meanings for sense bytes 0–7.

3.4.1. Sense Data Byte 0

When the bits of sense data byte 0 are set they have the following significance.

- Bit 0 – Command Reject

  Command Reject sense generates a Format 0 message which details the invalidity of the command. Command Reject indicates a programming error. Command Reject may appear with:

  – Write Inhibit (byte 1, bit 6). The disc unit READ ONLY switch (on 8430/8433 discs only) is enabled.

  – File Protect (byte 1, bit 5). The Write operation attempted is prohibited by the file mask.

- Bit 1 – Intervention Required

  Intervention Required sense indicates that the disc unit is either logically or electrically off line.

- Bit 2 – Bus Out Parity

  Bus Out Parity sense indicates that a parity check occurred on a command or data byte transfer from the channel.

- Bit 3 – Equipment Check

  Equipment Check sense produces a Format 0, 1, 2, or 3 Sense which details the error. Equipment Check with byte 1, bit 0 (Permanent Error) implies that the SCU has attempted recovery and failed.

SENSE BYTE NO.

	0	1	2	3	4	5	6 (8405/8430)	6 (8433)	7
BIT 0	COMMAND REJECT	PERMANENT ERROR	0	RESTART COMMAND	PHYSICAL ID	CYLINDER	REVERSE	0	FORMAT DECODE BITS 0-3 DECODES 7-E NOT USED F = INLINE SENSE DATA
BIT 1	INTERVENTION REQUIRED	INVALID TRACK FORMAT	CORRECTABLE				CYL (256)	CYL (512)	
BIT 2	BUS OUT PARITY	END OF CYLINDER	0				HIGH DIFF	CYL (255)	
BIT 3	EQUIPMENT CHECK	FORMAT 7 SENSE	0				HEAD		
BIT 4	DATA CHECK	NO RECORD FOUND	0						MESSAGE DECODE BITS 4-7
BIT 5	OVERRUN	FILE PROTECTED	0						
BIT 6	TRACK CONDITION CHECK	WRITE INHIBITED	0						
BIT 7	SEEK CHECK	OPERATION INCOMPLETE							

NOTE: BYTES 0 THRU 6 ARE THE SAME FOR ALL FORMATS.

MESSAGE DECODE	FORMAT 0 MESSAGE ONLY	FORMAT 1 DEVICE ERRORS	FORMAT 2 C. U. ERRORS	FORMAT 3 SELECTIVE RESET	FORMAT 4 ECC UNCORRECTABLE	FORMAT 5 ECC CORRECTABLE	FORMAT 6 USAGE/ ERROR COUNT	FORMAT 7 COMMAND RETRY	FORMAT F INLINE SENSE DATA
0	NO MESSAGE	NO MESSAGE	NO MESSAGE	NO MESSAGE	HA FIELD ECC UNCORRECTABLE	HA FIELD CORRECTABLE		HA FIELD UNCORRECTABLE	
1	INVALID COMMAND	SET SECTOR ERROR	ECC P1 OR P3 COMPARE ERROR		COUNT FIELD ECC UNCORRECTABLE	COUNT FIELD CORRECTABLE		COUNT FIELD UNCORRECTABLE	
2	INVALID SEQUENCE	ABNORMAL INTERRUPT FROM DRIVE	ECC P2 COMPARE ERROR		KEY FIELD ECC UNCORRECTABLE	KEY FIELD CORRECTABLE		KEY FIELD UNCORRECTABLE	
3	CCW COUNT LESS THAN REQUIRED				DATA FIELD ECC UNCORRECTABLE	DATA FIELD CORRECTABLE		DATA FIELD UNCORRECTABLE	
4	DATA VALUE NOT AS REQUIRED	NO WRITE CURRENT SENSE			HA FIELD NO SYNC BYTE FOUND			HA FIELD NO SYNC BYTE FOUND	
5	UNUSED	UNUSED			COUNT FIELD NO SYNC BYTE FOUND			COUNT FIELD NO SYNC BYTE FOUND	
6	CHANNEL DISCONTINUED RETRY	SET CYLINDER ERROR	UNUSED	UNUSED	KEY FIELD NO SYNC BYTE FOUND			KEY FIELD NO SYNC BYTE FOUND	
7	CHANNEL RETURNED INCORRECT RETRY CMD	SET HEAD ERROR			DATA FIELD NO SYNC BYTE FOUND			DATA FIELD NO SYNC BYTE FOUND	
8	UNUSED	SET DIFFERENCE ERROR			UNUSED	UNUSED		UNUSED	
9	UNUSED	FILE STATUS NOT AS EXPECTED			AM DETECTION FAILURE ON RETRY			AM DETECT FAILURE ON RETRY	
A	UNUSED	SEEK ERROR						UNUSED	
B	IMPROPER ALTERNATE TRACK POINTER	SEEK INCOMPLETE ON RETRY			UNUSED			UNUSED	
C	SERDES MALFUNCTION NO ST4's	NO INTERRUPT FROM DRIVE						HA FIELD CORRECTABLE	
D	UNUSED	UNUSED						COUNT FIELD CORRECTABLE	
E	UNUSED	RSV/REL CHECK						UNUSED	
F	RETRY BYTE COUNTER OR SECTOR VALUE INCORRECT	SHORT SELECT TIM OUT						UNUSED	

FORMAT 7 DISCONNECTED COMMAND RETRY

	BYTE 0	BYTE 1	BYTE 2	BYTE 3	BYTE 4	BYTE 5	BYTE 6 (8405/8430)	BYTE 6 (8433)	BYTE 7
P	0	X'10'	FILE MASK (INTERNAL FORM)	COMMAND CODE (ERROR)	PHYSICAL DRIVE ID	CYLINDER	8405/8430	8433	FORMAT DECODE BITS 0-7 7
1							REVERSE		
2							CYL (256)	CYL (512)	
3							HIGH DIFF	CYL (256)	
4							HEAD BITS 3-7		MESSAGE DECODE BITS 4-7
5									
6									
7									

Figure 3–2. Sense Data Bytes 0–7

- Bit 4 – Data Check

  The processing of ECC bytes following a field indicates a Data Error. When posted with byte 1, bit 0 (Permanent Error) the Data Check is uncorrectable and the message number in byte 7 indicates the nature of the error. There is no further recovery to be performed.

  When posted with byte 2, bit 1 (Correctable) it indicates a correctable data error in the data field or key field. Sense bytes 15 through 22 contain sufficient information to correct the error in storage.

  There is one instance where Data Check can be posted on a Write command. If an update write is performed on an overflow record in which the second or subsequent segment of the record has a data error in the home address or count field, Data Check and Correctable are posted. The data error can be serviced in the normal manner (the error pattern is zero).

- Bit 5 – Overrun

  The channel was late, in relation to the disc unit, to present a command or data byte or late in accepting a data byte. The Overrun sense may appear with byte 1, bit 0 (Permanent Error) in which case, recovery attempts have been made by the SCU.

- Bit 6 – Track Condition Check

  Track Condition Check sense indicates that a defective track or, in some situations, an alternate track has been encountered. Sense bytes 5–6 contain the defective address if a defective track is encountered or contains the alternate address if an alternate is detected.

  A defective track is detected for any Read (including SCU internal reads) or Search command other than Read Home Address, Search Home Address or Read Record Zero commands.

  The flag byte for HA and RO indicates alternative/defective track status.

  An alternate track is detected only for operations which cause track switching (multi-track command and overflow record processing) and certain control commands (e.g. Seek Head) while positioned on an alternate track.

- Bit 7 – Seek Check

  Seek Check indicates a Seek Incomplete or Seek Check has occurred. Seek Incomplete is detected as a time out and Seek Check is detected through a mismatch of the physical address bytes (before the HA and each count field) and the nominal Seek Address.

3.4.2. Sense Data Byte 1

When the bits of sense data byte 1 are set they have the following significance:

- Bit 0 – Permanent Error

  Permanent Error sense indicates the SCU internal error recovery facilities are exhausted. Permanent Error is always a modifier bit for other error indications.

- Bit 1 – Invalid Track Format

  Invalid Track Format sense indicates that a Write command has been specified to write past index. Invalid Track Format indicates a programming error.

- Bit 2 – End of Cylinder

  End of Cylinder sense indicates that a multi-track read or search or an overflow record operation has attempted head switching beyond the highest head address. End of Cylinder indicates a programming error.

- Bit 3 – Format 7 Sense

  Format 7 Sense indicates that a Data Check has occurred and the SCU is performing error recovery. Sense byte 7 indicates the type of Data Check which occurred.

  Format 7 Sense is part of an implementation of Disconnected Command Retry in which the SCU programs its next recovery operation and passes the information to the channel in the form of sense.

  If the system chooses to perform the recovery operation, the 24 bytes of Format 7 Sense must be given back to the SCU as the argument of the Retry Restart command. When the SCU has exhausted its recovery options (27 retries at varying head offsets), Permanent Error and Data Check are set. Format 7 Sense information is incompatible with any other format.

- Bit 4 – No Record Found

  No Record Found sense is posted after index has been encountered twice in the same chain without:

  – A read of the home address or data field

  – A Control, Write or Sense operation

  No Record Found indicates a programming error.

- Bit 5 – File Protect

  A Seek File Mask violation has been detected. An invalid Seek or implied Seek (multi-track or overflow) has been attempted. When posted with Command Reject (byte 0, bit 0), File Protect indicates an invalid Write operation has been attempted. File Protect indicates a programming error.

- Bit 6 – Write Inhibited

  Write Inhibited sense modifies Command Reject and indicates a Write was attempted on a disc unit where the READ ONLY switch (on the 8430/8433 discs only) is set.

- Bit 7 – Operation Incomplete

  Operation Incomplete sense indicates that an overflow record operation terminated prematurely. Operation Incomplete may occur alone or with:

  – Data Check

  – File Protect

  – End of Cylinder

  – Track Condition Check

  If Operation Incomplete occurs with one of the above, the other condition should be processed first and the Operation Incomplete processed just prior to restarting the command chain. Sense byte 3 contains the proper restart command code.

3.4.3. Sense Data Byte 2

Sense data byte 2 contains the Correctable sense bit (bit 1). Correctable sense is a modifier for the Data Check bit and indicates that the Data Check bit is correctable.

3.4.4. Sense Data Byte 3

Sense data byte 3 contains the Restart command for an Operation Incomplete (byte 1, bit 7) and is valid only in the Operation Incomplete and Format 7 Sense cases.

3.4.5. Sense Data Byte 4

Sense data byte 4 identifies the physical (hardwired) disc unit address. (Refer to 2.4.1.5.) The byte is formatted in the following way:

BIT	0	1	2	3	4	5	6	7
MEANING	0		3 OF 6 CODE					

Bit 0: ALWAYS READ

Bit 1: STRING NUMBER
(0 = FIRST STRING OF 8 DISC UNITS
1 = SECOND STRING OF UP TO 8 DISC UNITS)

3.4.6. Sense Data Byte 5

Sense data byte 5 contains the low order cylinder (i.e., C_2 of $C_1C_2H_1H_2$) currently addressed.

3.4.7. Sense Data Byte 6

Sense data byte 6 completes the physical seek address where the access arm is positioned. Bytes five and six are useful for error recovery restart command chain. The bit format of byte 6 is shown in Figure 3–2.

3.4.8. Sense Format 0 (Bytes 8–23) – Programming and SCU Errors

The principal significance of sense format 0 is to communicate programming or SCU errors. The message number in byte 7 points to the type of error. Byte 8 through 21 are zero. Byte 22 contains the sector decrement field and byte 23 contains the statistical integrator field.

3.4.9. Sense Format 1 (Bytes 8–23) – Disc Unit Checks

Sense format 1 transmits to the system information on a particular disc unit checks. The significance of the format 1 bytes is shown in Figure 3–3.

FORMAT 1 DEVICE ERRORS

Bit	BYTE 8 FILE STATUS	BYTE 9	BYTE 10 CONDITION AT UNSAFE	BYTE 11 SERVO UNSAFE	BYTE 12 R/W UNSAFE	BYTE 13 FC REG	BYTE 14 FB REG	BYTE 15 FT REG	BYTE 16	BYTE 17	BYTE 18	BYTE 19	BYTE 20	BYTE 21 CUDI(IS REG)	BYTE 22	BYTE 23
P																
0	INDEX ERROR	ZERO	LINEAR MODE		UNSAFE	CONTENTS OF CUDI BUS OUT	CONTENTS OF CUDI BUS IN	MODULE SELECT GATE	ZERO	ZERO	ZERO	ZERO	ZERO	DRIVE SELECTION ERROR	SECTOR DECRE-MENT	STAT INTE GRATOR
1	OFFSET ACTIVE		FIRST CYLINDER & LINEAR MODE					TAG GATE						TAG NOT VALID AT DEVICE		
2	SEEK INCOMPLETE		ACCESS READY		R/W READY UNSAFE			ENABLE TAG VALID CHECK						DEVICE CHECK		
3	SEEK/FORMAT COMPLETE		ODD CYLINDER											FILE BUS OUT PARITY ERROR		
4	ON LINE		DRIVE TO INNER GUARD BAND= 1 (NOT EVEN CYLINDER)= 0	ANY UNSAFE EXCEPT R/W	HEADS UNSAFE			DECODE 8 BIT						FILE BUS IN PARITY ERROR		
5	PACK CHANGE		TO DRIVE FORWARD 20 IPS= 1 (NOT DIFFERENCE = 1)= 0	PACK SPEED UNSAFE	PLO UNSAFE			DECODE 4 BIT						TAG BUS OUT PARITY ERROR		
6	BUSY		DRIVE FORWARD 5 IPS= 1 (NOT DIFFERENCE=0) =0	VELOCITY UNSAFE	DC WRITE UNSAFE			DECODE 2 BIT								
7	RECORD SEARCH IN PROGRESS		HEAD LOAD	30 VOLTS DC UNSAFE	AC WRITE UNSAFE			DECODE 1 BIT								

FORMAT 2 CONTROL UNIT ERRORS

Bit	BYTE 8 (CE REG) CONTROL CHECK	BYTE 9 SERDES (SE REG)	BYTE 10 ECC (EE REG)	BYTE 11	BYTE 12	BYTE 13 FC REG	BYTE 14 FB REG	BYTE 15 FT REG	BYTE 16	BYTE 17	BYTE 18	BYTE 19	BYTE 20	BYTE 21 CUDI (IS REG)	BYTE 22	BYTE 23
P																
0	CHANNEL BUS IN PARITY ERROR A/B	CUDI UNSAFE	NO. 1'S FROM SERDES SINCE LAST ECC RESET	ZERO	ZERO	CONTENTS OF CUDI BUS OUT	CONTENTS OF CUDI BUS IN	MODULE SELECT GATE	ZERO	ZERO	ZERO	ZERO	ZERO	DRIVE SELECTION ERROR	SECTOR DECRE-MENT	STAT INTE-GRATOR
1	CHANNEL A INTERFACE CHECK	SERDES WRITE PARITY CHECK	P0 OR WRITE ERROR					TAG GATE						TAG NOT VALID AT DEVICE		
2	CHANNEL B INTERFACE CHECK	SERDES READ PARITY CHECK	P1 OR P3 ERROR					ENABLE TAG VALID CHECK						DEVICE CHECK		
3	DATA TRANSFER CHECK	BIT RING	P2 ERROR											FILE BUS OUT PARITY ERROR		
4	SERDES, CUDI, ECC CHECK		P1 ERROR					DECODE 8 BIT						FILE BUS IN PARITY ERROR		
5	CHANNEL C INTERFACE CHECK	ECC ERROR	P3 ERROR					DECODE 4 BIT						TAG BUS OUT PARITY ERROR		
6	CHANNEL D INTERFACE CHECK							DECODE 2 BIT								
7								DECODE 1 BIT								

FORMAT 3 SELECTIVE RESET

Bit	BYTE 8 FAILING ADDRESS (AD REG)	BYTE 9	BYTE 10 C.U. ERROR 1 (ET REG)	BYTE 11	BYTE 12	BYTE 13 IG REG	BYTE 14 IF REG	BYTE 15	BYTE 16	BYTE 17	BYTE 18	BYTE 19	BYTE 20	BYTE 21	BYTE 22	BYTE 23
P																
0		MEMORY ADDRESS BIT 5	MICRO PROGRAM FORCED ERROR FLAG	ZERO	ZERO	WRITE LATCH	LONG SELECT	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO
1	MEMORY ADDRESS PARITY	MEMORY ADDRESS BIT 6	CLOCK ERROR			READ LATCH	FREEZE DATA TRANSFER									
2		MEMORY ADDRESS BIT 7	PC, CONTROL FIELD PARITY ERROR			LAST BYTE REQUEST	DATA SEPARATOR ENABLE									
3	MEMORY ADDRESS BIT 0	MEMORY ADDRESS BIT 8	PA, MEMORY ADDRESS PARITY ERROR			OPERATIONAL IN										
4	MEMORY ADDRESS BIT 1	MEMORY ADDRESS BIT 9	A–BUS PARITY ERROR			ADDRESS IN	BLOCK SWITCH TO CHANNEL A									
5	MEMORY ADDRESS BIT 2	MEMORY ADDRESS BIT 10	B–BUS PARITY ERROR			STATUS IN	BLOCK SWITCH TO CHANNEL B									
6	MEMORY ADDRESS BIT 3	MEMORY ADDRESS BIT 11	ALU COMPARE ERROR				BLOCK SWITCH TO CHANNEL C									
7	MEMORY ADDRESS BIT 4	MEMORY ADDRESS BIT 12	PN, CONTROL MEMORY ADDRESS PARITY ERROR				BLOCK SWITCH TO CHANNEL D									

FORMAT 4 ECC UNCORRECTABLE

Bit	BYTE 8 CYLINDER	BYTE 9	BYTE 10 HEAD	BYTE 11	BYTE 12 RECORD	BYTE 13 SECTOR	BYTE 14 ACCESS OFFSET	BYTE 15 RETRY COUNT	BYTE 16 SOURCE PHYSICAL ADDRESS	BYTE 17	BYTE 18	BYTE 19	BYTE 20	BYTE 21	BYTE 22	BYTE 23
P																
0	HIGH ORDER BYTE	LOW ORDER BYTE	HIGH ORDER BYTE	LOW ORDER BYTE	RECORD NUMBER OF RECORD IN ERROR	SECTOR NUMBER OF START OF RECORD IN ERROR		NUMBER OF RETRYS REQUIRED TO RECOVER FROM THE ERROR	READ FROM ID BYTE WRITTEN BEFORE EACH RECORD, IDENTIFIES CU AND DRIVE THAT WROTE THE RECORD. BITS 0–1 CODE FOR CU BITS 2–7 DRIVE 3 FO 6 CODE A= 111000 E= 011100 B= 110001 F= 010101 C= 101010 G= 001110 D= 100011 H= 000111	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO	ZERO
1																
2							800 µINCHES									
3							400 µINCHES									
4							200 µINCHES									
5							100 µINCHES									
6							50 µINCHES									
7	(LAST SEEK ADDRESS)		(LAST SEEK ADDRESS)				25 µINCHES									

FORMAT 5 ECC CORRECTABLE

Bit	BYTE 8 CYLINDER	BYTE 9	BYTE 10 HEAD	BYTE 11	BYTE 12 RECORD	BYTE 13 SECTOR	BYTE 14 ACCESS OFFSET	BYTE 15 RESTART DISPLACEMENT	BYTE 16	BYTE 17	BYTE 18 ERROR DISPLACEMENT	BYTE 19	BYTE 20	BYTE 21 PATTERN	BYTE 22	BYTE 23 CHANNEL TRUNCATION
P																
0	HIGH ORDER BYTE	LOW ORDER BYTE	HIGH ORDER BYTE	LOW ORDER BYTE	RECORD NUMBER OF RECORD IN ERROR	SECTOR NUMBER OF START OF RECORD IN ERROR		HIGH ORDER		LOW ORDER	HIGH ORDER	LOW ORDER	HIGH ORDER		LOW ORDER	
1																
2							800 µINCHES									
3							400 µINCHES									
4							200 µINCHES									
5							100 µINCHES	COUNT OF TOTAL NUMBER OF BYTES READ IN RECORD. CAN INCLUDE OVERFLOW SEGMENTS.			COUNT OF NUMBER OF BYTES FROM FIRST ECC BYTE TO START OF DATA IN ERROR.		ECC GENERATED PATTERN TO CORRECT DATA IN ERROR. INCLUDES ALL THREE BYTES.			
6							50 µINCHES									
7	(LAST SEEK ADDRESS)		(LAST SEEK ADDRESS)				25 µINCHES									

FORMAT F INLINE SENSE DATA

Bit	BYTE 8 FILE STATUS 1	BYTE 9 STATUS 2	BYTE 10 SERVO STATUS	BYTE 11 SERVO UNSAFE	BYTE 12 R/W UNSAFE	BYTE 13 FC REG	BYTE 14 FB REG	BYTE 15 FT REG	BYTE 16 CAR	BYTE 17 HAR	BYTE 18 DIFF	BYTE 19 SAR	BYTE 20	BYTE 21 CUDI(IS REG)	BYTE 22	BYTE 23
P																
0	INDEX ERROR	INDEX ERROR	LINEAR MODE			CONTENTS OF CUDI BUS OUT	CONTENTS OF CUDI BUS IN	MODULE SELECT GATE	NOTE: In the SEEK INCOMPLETE case, CAR, HAR, & DIFF are the values loaded prior to the seek. Bytes 5 & 6 are the seek address prior to the seeking. In the SEEK COMPLETE case, bytes 5 & 6 are the same as bytes 10 & 11.				CORRECT REGISTER CONTENTS OR SECTOR UNDER TEST (BYTE 16, 17, 18, OR 19)	DRIVE SELECTION ERROR	ZERO	ZERO
1	OFFSET ACTIVE	OFFSET ACTIVE	FIRST CYLINDER & LINEAR MODE					TAG GATE						TAG NOT VALID AT DEVICE		
2	SEEK INCOMPLETE		ACCESS READY		R/W READY UNSAFE			ENABLE TAG VALID CHECK						DEVICE CHECK		
3	SEEK/FORMAT COMPLETE	READ ONLY	ODD CYLINDER		BUS OUT INVALID									FILE BUS OUT PARITY ERROR		
4	ON LINE	R/W READY	DRIVE TO INNER GUARD BAND= 1 (NOT EVEN CYLINDER)= 0	ANY UNSAFE EXCEPT R/W	HEADS UNSAFE			DECODE 8 BIT TAG						FILE BUS IN PARITY ERROR		
5	PACK CHANGE	INDEX	TO DRIVE FORWARD 20 IPS= 1 (NOT DIFFERENCE = 1)= 0	PACK SPEED UNSAFE	PLO UNSAFE			DECODE 4 BIT TAG						TAG BUS OUT PARITY ERROR		
6	BUSY	END OF CYL	DRIVE FORWARD 5 IPS= 1 (NOT DIFFERENCE=0) =0	VELOCITY UNSAFE	DC WRITE UNSAFE			DECODE 2 BIT TAG								
7	RECORD SEARCH IN PROGRESS	WRITE CUR SENSE	HEAD LOAD	30 VOLTS DC UNSAFE	AC WRITE UNSAFE			DECODE 1 BIT TAG								

Figure 3–3. Sense Byte Detail

3.4.10. Sense Format 2 (Bytes 8–23) – SCU Checks

Sense format 2 transmits to the system information on a particular SCU check. The significance of the format 2 bytes is shown in Figure 3–3.

3.4.11. Sense Format 3 (Bytes 8–23) – Selective Reset

The state of the SCU just prior to a Selective Reset is transmitted through sense format 3. The SCU enters a "hard wait state" if it detects an internal catastrophic error condition. As the SCU enters the "hard wait state", it notifies the processor that a Selective Reset must be performed to make the SCU operational again.

After receiving a Selective Reset, the SCU is operational and it conditions itself such that the next I/O sequence issued to any device address on that controller receives an immediate Unit Check. The sense information indicates Equipment Check.

The significance of the format 3 bytes is shown in Figure 3–3.

3.4.12. Sense Format 4 (Bytes 8–23) – ECC Uncorrectable Data Error

The details of an uncorrectable data error are transmitted to the system in sense format 4. The significance of the format 4 bytes is shown in Figure 3–3.

3.4.13. Sense Format 5 (Bytes 8–23) – ECC Correctable Data Errors

Information on data error correction and I/O program continuation is transmitted by sense format 5. The significance of the format 5 bytes is shown in Figure 3–3.

3.4.13.1. Data Error Correction

Data errors in home address, count fields, and key fields are corrected internally by the SCU. However, when the command is a Read, the key field errors are not corrected by the SCU.

Errors in data fields must be corrected by the processor.

Sense bytes 20 through 22 contain an error pattern which must be aligned and exclusive ORed to the erroneous data in storage. Sense bytes 20 through 22 may be considered to be a single 24-bit quantity which must be aligned bit for bit to the erroneous 24-bit quantity. Note that with scatter reads, the erroneous 24-bit reads may not be contiguous in storage.

Sense bytes 18 and 19 are a 16-bit unsigned number giving the number of bytes between the first byte in error and the end of the field in error.

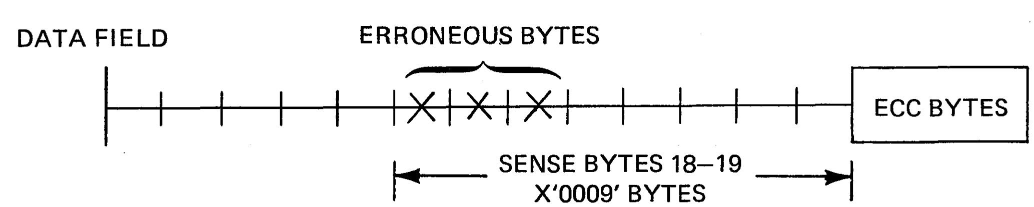

Sense bytes 15–17 are 24-bit unsigned numbers giving the number of bytes between the first byte transferred by the particular read command and the end of the field in error.

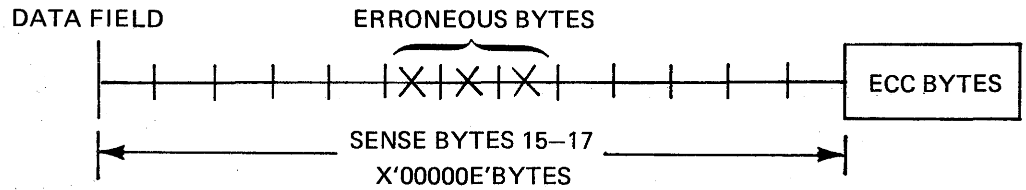

Note that sense bytes 15–17 are partially determined by the command. For the same Data Field Error, bytes 15–17 will vary (compensate) for Read Count, Key and Data, Read Key and Data, etc. The example assumes a Read Data command.

Subtracting sense bytes 18 and 19 from bytes 15–17 yields the forward displacement relative to the first byte transferred by the particular Read command.

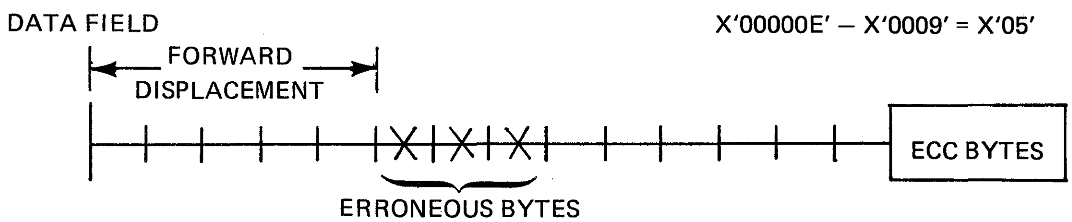

It remains only to exclusive OR sense bytes 20–22 to data bytes 5–7 respectively after determining how bytes 5–7 are distributed in storage.

NOTE:

In case the error occurred partially within the ECC bytes themselves, the SCU clears the error pattern corresponding to the erroneous ECC bytes.

3.4.13.2. I/O Program Continuation

If the data error occurred before the command sequence ended, a command sequence must be developed to bring the remainder of the interrupted program to completion.

- Sense bytes 5 and 6 can be used to develop the proper continuation seek address
- Sense byte 13 contains the sector number of the record in error
- Sense bytes 8–12 contain the CCHHR identifier field of the record in error
- Sense byte 23, bit 7 if on, indicates that the channel truncated during the read

3.4.14. Sense Format 7

Disconnected Command Retry sense provides the argument to the Retry Restart command. See Figure 3–2 for detail byte meanings.

3.4.15. Sense Format F (Bytes 8–23) – Inline Sense

Format F sense is never transmitted to the system. It is available only at the SCU CE panel and only during the running of inline diagnostics. The significance of the Format F bytes is shown in Figure 3–3.

3.5. PROGRAMMING CONSIDERATIONS

Programming and operating conventions pertaining to disc unit fault indications, overflow record processing, and logical end-of-file processing are described in the following paragraphs.

3.5.1. Disc Unit Fault Indications

The following conditions cause an UNSAFE (8430/8433 discs) or DEVICE CHECK (8405 FHD) indication on the disc unit control/indicator panel and result in a lockout of the addressed disc unit:

- More than one read/write head selected simultaneously
- Read and write controls selected simultaneously
- Read and erase controls selected simultaneously
- AC power loss
- DC voltage variation beyond specified limits
- Write current flowing without write gate up
- Head selected without seek complete
- Write gate up without ac power
- Incorrect disc pack rotational speed

3.5.2. Overflow Records

An overflow record is a logical record that exceeds the capacity of one track. The part of an overflow record that is on one track is referred to as a record segment. Each segment contains a count field, key field (optional), and a data field. The key length and data length specified in the count field (KL, DL subfields) refer only to that segment, not the entire overflow record.

To format an overflow record, all segments but the last are written with a Write Special Count, Key and Data command (01_{16}). The Write Special Count, Key and Data command causes a 1 to be written in bit four of the flag byte indicating an overflow record. The SCU does not automatically head switch while an overflow record is being formatted. All head seeking must be done by the formatting program. The last segment of an overflow record must be written by a Write Count, Key and Data command (10_{16}).

All segments of an overflow record, except the first, must immediately follow the Track Descriptor Record (Record Zero). All segments of an overflow record, except the last, must be the last physical record on their respective tracks.

The following commands may be used to read or update previously written overflow records:

- Write Data
- Write Key and Data
- Read Data
- Read Key and Data
- Read Count Key and Data

When an overflow record is detected, the SCU processes the first segment and waits for the index mark. When the index mark is detected, the head address is incremented by 1. The home address of the new track is read to verify successful switching.

The SCU clocks (reads internally) the Track Descriptor Record (Record Zero) to continue processing in the data area until a segment without a 1 in flag byte bit position 1 is detected. The operation concludes when the data area of the segment is processed.

3.5.3. End-Of-File

An end-of-file record is used to indicate the end of a logical group of records. To write an end-of-file record, a Write Count, Key and Data command is used. The data length (DL) subfield of the count field must equal zero. The DL=0 condition causes the SCU to write one byte of zeros followed by the ECC bytes.

When the count field (DL=0) of an end-of-file record is read, Unit Exception status is posted by the SCU.

3.6. ERROR RECOVERY

System-Subsystem error recovery techniques and capabilities are discussed in the following paragraphs.

3.6.1. Error Correction Code

Error detection and correction is aided by a mathematically derived error correction code (ECC). The implementation of this code into hardware has given the SCU the ability to detect single burst errors of 22 bits or less and correct single burst errors of 11 bits or less.

During a write operation, data is transferred from the channel, serialized and sent to the disc unit. As the serial data is being transferred to the disc unit, it also is fed through the ECC circuitry. The ECC circuit generates, as a function of the input serial data, a unique code for error detection and correction. This code (7 bytes) is written on the disc following the data field that generated it.

During a read operation, data is transferred from the disc unit and fed through the ECC circuits. The seven ECC bytes are also read and fed to the ECC circuit. The properties of the error correction code are such that both the data and ECC bytes must have been read correctly or an error is detected by the ECC circuit.

3.6.2. Transparent Command Retry

Transparent command retry allows recovery to take place without system Software intervention. The basic properties of transparent command retry are:

- Transparent command retry must be allowed (file mask bit 6=1). If no Set File Mask command is issued, transparent command retry is not allowed. (Exceptions: The Read IPL command, recovery from offset, and index encountered while clocking data are all instances where transparent command retry is always allowed.)

- Transparent command retry is invoked only when the SCU determines that the error is recoverable in one rotation.

- Transparent command retry remains active only for the duration of the first data field processed by any I/O Command Sequence, and becomes disallowed for the remainder of the sequence.

3.6.2.1. Index-Encountered-While-Clocking-Data

Index-encountered-while-clocking-data occurs due to a possible skew of the read/write heads in relation to the servo head.

No status is presented by the SCU. The SCU, MSA and IOU remain connected through the retry. The retry does not require the command to be reissued.

The recovery requires one rotation (16.7 milliseconds).

Transparent command retry is always allowed for this condition.

3.6.2.2. Recovery-From-Offset

Recovery-from-offset occurs when an I/O Command Sequence is initiated to a disc unit in which the heads have been offset due to a previous error recovery operation.

No status is presented by the SCU. The SCU, MSA and IOU remain connected through the retry. The retry does not require the command to be reissued.

The recovery requires one rotation (16.7 milliseconds).

Transparent command retry is always allowed for this condition.

3.6.2.3. Correctable-Data-Errors-While-Searching

Correctable-data-errors-while-searching occurs when a correctable data error is detected in a home address, count or key field being searched.

Channel End is presented to the MSA and followed later by Device End. The MSA will reinitiate the Search (since there was no Status Modifier). The status is not propagated by the MSA. The SCU, MSA, and IOU remain connected through the retry.

The recovery requires one rotation (16.7 milliseconds).

3.6.2.4. Correctable-Data-Error-in-a-Clocked-Home-Address-or-Count-Field

The condition occurs due to detecting a data error in a home address or a count field which is being read internally by the SCU, but the field is not being transferred to the MSA.

No status is presented by the SCU. The SCU, MSA and IOU remain connected through retry. The retry does not require the command to be reissued.

3.6.3. Disconnected Command Retry

Disconnected Command Retry (DCR) is a data error recovery tool in which the SCU and disc unit can be reconditioned to the context of an error long after the error has occurred. In the interim period between the occurrence of the error and the subsequent retry, the disc unit, SCU, MSA and IOU are free for other operations. The recovery can therefore be interleaved with other operations in a discontinued or disconnected form.

3.6.3.1. Status Presentation

Channel End, Device End and Unit Check status bits are generated at the occurrence of any Disconnected Command Retry recovery condition. A subsequent Sense command will show sense byte 1, bit 3 set to one indicating Format 7 Command Retry Sense. After the sense command has been issued, the subsystem is usable for operations other than the retry.

3.6.3.2. Recovery Restart Attempt

To attempt recovery the following actions are required:

- Sense bytes 0 through 23, obtained from the Format 7 Sense, must be given to the SCU as the argument of the Retry Restart command $3B_{16}$ (73_8).

- Chained to the Retry Restart command must be the command upon which the original error occurred. Note that under the conditions of data chaining, the first data chained command in the sequence which failed is the command which must be restarted.

3.6.3.3. Error Types Serviced by Disconnected Command Retry

Disconnected Command Retry may be used to recover from:

- Uncorrectable errors in any field

- Correctable data errors in home address or count fields under Read Home Address, Read Count or Read Count Key and Data commands

- Correctable data errors in home address or count fields being clocked (read internally by the SCU but not transferred to the MSA) when transparent command retry is disabled

- Correctable data errors in home address and count fields under Search commands when transparent command retry is disabled. Note that correctable key field errors always result in sense format 52_{16} errors when transparent command retry is disabled

3.6.3.4. Disconnected Command Retry Usage

Disconnected Command Retry usage results in one of four cases:

- A reoccurence of the same or different error (i.e., a successive Disconnected Command Retry situation)
- A retry is the 27th and final recovery attempt and fails with Channel End, Device End and Unit Check status. A subsequent Sense command will show Permanent Error posted (sense byte 1, bit 0)
- The retry is successful and the I/O Command Sequence terminates normally
- The retry is successful resulting in a correctable data error in a data field or possible a correctable data error in a key field (when the command is a Read)

3.6.3.5. Disconnected Command Retry Notes

If the drive incurring the error has the dual-access feature and the physical address plugs for each of the accesses are identical, the retry may be attempted from the alternate SCU. If the logical address plug of the dual-access 8430/8433 has been changed on the disc unit that had the error in the time between the error and the recovery attempt, the recovery sequence should be issued to the new logical address.

When performing DCR recovery on a data check, intermixed with other I/O requests to the disc unit that incurred the error, the other I/O requests may encounter a number of recovery-from-offset conditions.

3.7. TIMING

The basic timing components for the subsystem are accessor positioning time and disc latency time.

3.7.1. Accessor Positioning Time

NOTE:

Accessor positioning time applies to the 8430/8433 discs only.

Positioning time is defined as the time interval from the beginning of a Seek Start (from zero offset) to the generation of a Seek Complete by the disc. The minimum and maximum positioning times are:

Head Positioning	Time (in milliseconds)	
	8430	8433
track-to-track (min.)	7	10
average	27	30
full travel (max.)	50	55

The minimum and maximum positioning time is the average of at least 1024 Seeks. The average positioning time is the sum of the time required to do all the possible combinations of Seeks divided by the number of such Seeks.

3.7.2. Disc Latency Time

Disc latency time is the time interval between the instant at which the SCU initiates a call for data and the start of the actual transfer of the data. The specific disc latency times are:

- ½ revolution (average) – 8.33 milliseconds
- Full revolution (maximum) – 16.67 milliseconds

3.8. MSA FUNCTIONS

The MSA receives functions from the processor via the function word. These functions can be executed by the MSA without issuing commands to the disc subsystem. A brief description of the buffered and unbuffered MSA functions is provided in the following paragraphs. A description of the MSA word formats is given in Appendix A.

NOTE:

Only the test and jump functions are used in both the buffered and unbuffered MSAs, the remaining functions are used only in the buffered MSA.

3.8.1. Test

The test function is used to return detailed status information to the processor. Once issued to the MSA, the test function returns the auxiliary status word by way of an external interrupt. To select the MSA, the DA field of the function word must be $F1_{16}$. The test function is to be used when a regular status word is returned indicating an MSA error (bit 16 is 1). The cause of the error and the address of the function in which the error was detected is contained in the auxiliary status word. The particular bits set in the status byte are not examined or altered by the MSA. The MSA disconnects from the disc subsystem once status transfer is completed. Status words are generated in response to the test function.

3.8.2. Jump

The jump function repeats a chain of functions or a section of the chain without interrupting the processor. When used, it must always be the last function in the chain.

3.8.3. Set Queue Limit

The set queue limit function allows the MSA to preset a queuing limit in the word buffer control. The set queue limit function overrides a hardware constant in the MSA that specifies the MSA must receive 64 words of output data before a command is issued to the disc subsystem. When an error is detected during the execution of this function, the normal status word will contain the E-bit set to one and the $F1_{16}$ in the device address field. The device status field will contain the command code of the erroneous function.

3.8.4. Set Partial Byte Transfer

The set partial byte transfer function permits the MSA to transfer a word to the processor containing a partial byte. Normally the partial byte is stripped by the MSA when operating in format C. The MSA will present the normal status word on completion of this function if it is a single function or the last function.

3.8.5. Reset Partial Byte Transfer

The reset partial byte transfer function resets the partial byte transfer to the normal mode of operation when partial byte in the last input word will be stripped by the MSA.

Appendix A. MSA Word Formats

A.1. GENERAL

The MSA accepts and transfers function words, status words, auxiliary status words, and data words. A brief description of these words is provided in the following paragraphs. This description is applicable to both the buffered and unbuffered MSA, except where noted otherwise. A complete description of these words and the operation of the MSA is contained in SPERRY UNIVAC 1100 Series Multi-Subsystem Adapter (MSA) Programmer Reference, UP–7890 (current version).

A.2. FUNCTION WORD

The basic function word contains two basic sections. The least significant 16 bits specifies the command code and the device address for the disc subsystem. The next 12-bit field (bits 16–27) is interpreted by the MSA and specifies various formatting and control functions internal to the MSA. Figure A–1 shows the detailed function word format. An explanation of the bit position of the function word is given in following paragraphs.

B	C* L R	NOT USED (MUST BE ZEROS)	P	S	R	C	M	X	Q	F	CC**	DA***
35	34	33 28	27 24	23	22	21	20	19	18	17 16	15 8	7 0

LEGEND:

B — Bootstrap
CLR — Soft I/O Clear
P — Parameter Count/Jump Address
S — Search Flag
R — Channel Reserve Flag
C — Command Chaining Flag
M — Multiple Function Flag
X — Translate
Q — Queuing Hold
F — Format Select
CC — Command Code
DA — Device Address

*Used in the buffered MSA only. In the unbuffered MSA this bit position must be zero.
**CC (bit positions 15 through 8) must be equal to 00_{16}. All other bits are ignored.
***DA (bit positions 7 through 0) must be equal to $F1_{16}$. All other bits are ignored.

Figure A–1. Function Word Format

A.2.1. Bootstrap (B)

The bootstrap field (B), in bit position 35, specifies bootstrap operation and occurs when power is being restored to the processor or MSA after a power shutdown or loss. The MSA automatically selects data format C (see A.5) for bootstrap operation.

A.2.2. Soft I/O Clear (CLR)

The soft I/O clear field (CLR), in bit position 34, (used in the buffered MSA only) clears the active channel when the I/O clear is detected on the inactive channel. The I/O clear is not detected on the active channel because the active channel gates data words with the EF or the OA.

NOTE:

1110 Processing System will clear on active channel.

A.2.3. Parameter Count/Jump Address (P)

The parameter count/jump address field (P), in bit positions 27 through 24, is used as an indication that output data is required on an output function if the search flag is not specified. When the search flag is used, the P field indicates the upper limit of parameter bytes to be stored by the MSA in the search identifier register (SIR) in the unbuffered MSA. (In the buffered MSA, the parameter bytes are transmitted by the MSA to the disc subsystem.) Normally, this count should be set to the number of bytes required by the disc subsystem. The following specifies the P field definitions as related to the search flag (S):

P Bits 27–24	S Bit 23	Interpretation
1XXX	0	Output data required for Control or Write command
0XXX	0	Output data not required for Control or Write command
0000 1100	1	Binary count indicating upper limit of bytes to be stored in SIR: 0–12
1101 1111	1	Indicates a binary count of 12 (maximum)

NOTE:

Xs in bit format equal 0 or 1.

On a jump function, the three least significant bits (bits 26, 25, and 24) of the P field are used to specify the internal address of the next function to be executed in the chain. The coding is binary, that is, 000 specified the first function, 001 the second, etc.

A.2.4. Search Flag (S)

The search flag field (S), in bit position 23, indicates that the specified command is to be reissued to the disc subsystem, depending on the status indication received from the disc subsystem in the buffered MSA, or using the SIR as the source of output data in the unbuffered MSA. When the first command is executed, parameter data is simultaneously transmitted to the SCU and loaded in the SIR in the unbuffered MSA. Subsequent commands use data previously stored in the SIR until the function is completed. In the buffered MSA, the parameter data stored in the output buffer is transmitted to the disc subsystem. The number of bytes transmitted to the SCU (in the buffered MSA) or to be stored in the SIR (in the unbuffered MSA) is specified by the P field (see A.2.3). The function is completed normally when a status byte consisting of the Status Modifier and Device End bits is returned by the disc subsystem. A returned error status which suppresses chaining will also terminate the function in which the search flag is specified.

If the number of bytes requested by the disc subsystem is greater than the count specified in the P field, the SIR is loaded with the count specified in the unbuffered MSA or only the specified count will be transmitted (buffered MSA) and data transfer to the disc subsystem terminates.

If the number of bytes requested is less than the P field count, the SIR is loaded with the number of bytes requested by the device (in the unbuffered MSA), and the P field is ignored. In the previous two cases, the commands continue to be issued and the function termination becomes dependent on the status received.

A.2.5. Channel Reserve Flag (R)

The channel reserve flag field (R), in bit position 22, specifies a particular mode of operation when the MSA is accessed by two processor channels. Once a function word, or chain of function words, has been stored in the MSA, the channel reserve flag controls the channel switching after operation is complete. Normally, when an operation is complete, function words are acted upon in order of channel priority; however, if R=1 in any function word, the channel connection is maintained after completion of the operation and until a function or chain of functions is successfully completed on that channel where the channel reserve flag is not specified.

A.2.6. Command Chaining Flag (C)

The command chaining flag field (C), in bit position 21, specifies that a group of functions is to be executed in a chain to a single device. Once the initial function word is stored in the MSA, this bit indicates that further function words are to be requested and stored in the MSA function buffer for subsequent execution. The MSA continues to store functions until C is reset to 0 or the MSA function buffer is full. Command chaining implies that the MSA stays connected to the disc subsystem until the command chain is broken or the last command in the chain is executed. A status is returned at the completion of the execution of the chain. The initial selection sequence is not initiated for the disc subsystem until all function words are stored.

A.2.7. Multiple Function Flag (M)

The multiple function flag field (M), in bit position 20, enables the MSA to execute a string of functions to different devices without storing those functions in the MSA. In this mode of operation, the MSA, upon completion of the first function, interprets the next full output word as a function word and executes the second function. The MSA keeps calling out functions until M=0, which defines the last function in a string.

A.2.8. Translate (X)

The translate field (X), in bit position 19, specifies a translation of output or input data in the unbuffered MSA, the buffered MSA contains only one translator. Up to two bidirectional translators can be installed in the same MSA module. Each processor I/O channel is assigned to one MSA translator or the other by way of patchcard wiring. Up to four I/O channels can be assigned to a single translator. The two available translators provide:

- Fieldata/ASCII translation
  - Processor Fieldata code (64 characters) to device ASCII code (64 characters)
  - Device ASCII code (96 characters) to processor Fieldata code (64 characters)
- Fieldata/EBCDIC translation
  - Processor Fieldata code (64 characters) to device EBCDIC code (64 characters)
  - Device EBCDIC code (64 characters) to processor Fieldata code (64 characters)

A.2.9. Queuing Hold (Q)

The queuing hold field (Q), in bit position 18, is used to format data. When the Q bit is set to 1, partially disassembled data on output commands is not discarded upon completion of the command. The first byte for the following command is the next byte of the partially disassembled word. The following command need not be an output type command because the data is saved until an output type command is issued.

The Q bit also can be used on input operation. If the input records do not fill a complete word upon termination of a command, the data is held until the input data from a subsequent command fills out that word. If the chain is terminated with a partial word filled, the MSA reverts to its normal mode of operation by loading the specified word into its input queue for transmission to the processor.

The Q bit is for use only with commands requiring output or input data transfer.

A.2.10. Format Selection (F)

The format selection field (F), in bit positions 17 and 16, specifies format A, B, or C of the word/byte conversion for data words. Each function can specify a different format; however, data words transmitted under one function can only be in one of the three formats. The following chart illustrates the format selections:

Bits		
17	16	Format
0	0	A – Quarter word
0	1	B – Packed 6-bit word
1	0 or 1	C – Packed 8-bit word

On bootstrap operations, format C is selected automatically. See A.5 for details of data formats.

A.2.11. Command Code (CC)

The command code field (CC), in bit positions 15 through 8, is transmitted to the disc subsystem as the command byte in the initial selection sequence. The general format interpreted by the MSA is:

Bit Locations 15 8		Operation	
0 0 0 0	0 0 0 0	Test	
X X X X*	X X 0 1	Write	Output Forward
X X X X	X X 1 1	Control	
X X X X	0 1 0 0	Sense	Input Forward
X X X X	X X 1 0	Read	
X X X X	1 1 0 0	Input Backward	
X X X X	1 0 0 0	MSA Command (jump)	

NOTE:

Xs is bit format equal 0 or 1.

A.2.12. Device Address (DA)

The device address byte (DA), in bit positions 7 through 0, specifies the SCU and disc unit. The DA format is:

0	1	2	3	4	5	6	7	Device Address Byte Bit Positions
SCU					Disc Unit			

The MSA is assigned device address $F1_{16}$ (1111 0001); SCU addresses must have byte bit 0 (bit position 7 of function word) = 0. A disc unit always has the same device address, but the SCU postion of the device address may vary for different I/O channels. Devices must be numbered consecutively from 0.

A.3. STATUS WORD

The 36-bit word is sent to the system at the completion of each function or chain of functions. In the unbuffered MSA, status information occupies the least significant 18 bits only; the most significant 18 bits are not used. In the buffered MSA, status information occupies bit positions 0 through 35. In both MSA types, the least significant 16 bits relate to the device address and device status; and bits 16 and 17 refer to error conditions detected in the MSA. The remaining bits, 18 through 35, are used only in the buffered MSA and refer to the abnormal byte count, residual word count, and error condition fields. The status word format for the buffered and unbuffered MSA is shown in Figure A–2. This status word applies only to functions addressed to the device and is called the normal status word to differentiate it from the auxiliary status word (see A.4). An explanation of the bit positions of the normal status word is given in the following paragraphs.

O	CON	ACP	RWC		M		ABC	E	DS		DA	
35	34	33	32	22	21	18	17	16	15	8	7	0

a. Buffered MSA Status Word Format

NOT USED (ANY COMBINATION OF 0 AND 1 BITS)		ABC	E	DS		DA	
35	18	17	16	15	8	7	0

b. Unbuffered MSA Status Word Format

LEGEND:

Bit 35 – Normal Status Word Indicator, always zero (Buffered MSA only)
CON – Contingency Error (Buffered MSA only)
ACP – Address Compare Error (Buffered MSA only)
RWC – Residual Word Count (Buffered MSA only)
M – Magnitude of Byte Count Error (Buffered MSA only)
ABC – Abnormal Byte Count
E – MSA Error
DS – Device Status

ATT	– Attention	Bit 15
SM	– Status Modifier	Bit 14
CUE	– Control Unit End	Bit 13
B	– Busy	Bit 12
CE	– Channel End	Bit 11
DE	– Device End	Bit 10
UC	– Unit Check	Bit 9
UE	– Unit Exception	Bit 8

DA – Device Address

Figure A–2. Normal Status Word Format

A.3.1. Normal Status Word Indicator

Bit position 35 is the normal status word indicator and is always zero. Bit 35 indicates that the status word transmitted to the processor is the normal status word.

A.3.2. Contingency Error (CON)

The contingency error field (CON), in bit position 34, is used only in the buffered MSA status word. The CON field indicates that the MSA attempted to reset contingency connection, but received error indication from the SCU, i.e. status byte from the disc subsystem contained bits other than Channel End and Device End.

A.3.3. Address Compare Error (ACP)

The address compare error (ACP), in bit position 33, is used only in the buffered MSA status word. The ACP field is the same as bit 13 of the auxiliary status word (see A.4). The ACP field indicates that the device address transmitted to the disc subsystem did not agree with the returned address by the disc subsystem.

A.3.4. Residual Word Count (RWC)

The residual word count field (RWC), in bit positions 32 through 22, is used only in the buffered MSA status word. The RWC field indicates the number of words remaining in the word output buffer upon termination of operation.

A.3.5. Magnitude of Byte Count Error (M)

The magnitude of byte count error field (M), in bit positions 21 through 18, is used only in the buffered MSA status word format. The M field indicates the position of the last significant information byte/character in the last word transmitted to or received from the disc subsystem.

A.3.6. Abnormal Byte Count (ABC)

The abnormal byte count (ABC), in bit position 17, indicates that the number of output bytes transmitted to the device does not correspond to an integral number of 36-bit output data words in a particular format selection. On input operations, ABC is set when the number of data bytes received from the device does not completely fill the last data word.

A.3.7. MSA Error (E)

The MSA error bit (E), in bit position 16, indicates that a programming or hardware error was detected in the MSA. The detailed information regarding the particular error is contained in the auxiliary status word. An error condition detected in the device does not set the E bit.

A.3.8. Device Status (DS)

Bit definitions for the device status field (DS), in bit positions 15 through 8, are given in 3.3.

A.3.9. Device Address (DA)

The device address field (DA), in bit positions 7 through 0, indicates the 8-bit address of the device to which the function or chain of functions was addressed. If, after the MSA disconnects itself from a device, the MSA receives a status byte from the device, the device address in the resultant status word shows the address of the device associated with that status byte.

A.4. AUXILIARY STATUS WORD

The auxiliary status word contains detailed information relating to error conditions detected in the MSA. The word must be requested by the processor by sending a test function command code with the MSA device address $F1_{16}$ (1111 0001). This test function causes the auxiliary status word to be made available with an external interrupt signal. The acceptance of the auxiliary status word signals successful completions of the test operation. The normal status word does not apply unless the test function is not successfully executed.

The axuiliary status word should be requested by way of the test function when the E bit is set in the normal status word or when other information regarding error conditions is required. Unlike the normal status word, the auxiliary status word remains stored in the MSA until a new function or chain of functions is initiated to the disc subsystem. The MSA test function does not clear the content of the auxiliary status word. The format for the auxiliary status word is shown in Figure A–3. An explanation of the bit positions of the auxiliary status word is given in the following paragraphs.

				INTERFACE CONTROL CHECKS								PROGRAM CHECKS				
1	NOT USED	SDIR	EDC	NOP	IPE	TCK	CLE	ACP	XCK	LA	SC	WPC	IS	IC	FA	M
35	34 – 24	23	22 – 18	17	16	15	14	13	12	11	10	9	8	7	6 – 4	3 – 0

LEGEND:

Bit 35	– Auxiliary Status Word Indicator, always 1 (Buffered MSA only)	
SDIR	– Service In/Data In Received (Buffered MSA only)	
EDC	– Error Detection Code (Buffered MSA only) (bit 22 through 7 for buffered)	
NOP	– Control Unit Nonoperational	Byte interface control checks
IPE	– Input Parity Error	Byte interface control checks
TCK	– Time Check	Byte interface control checks
CLE	– Control Line Error	Byte interface control checks
ACP	– Address Compare Error	Byte interface control checks
XCK	– Translate Check	
LA	– Late Acknowledge	
SC	– Stall Check	
WPC	– Word Parity Check	Program checks
IS	– Invalid Sequence	Program checks
IC	– Invalid Command	Program checks
FA	– Function Address	Program checks
M	– Magnitude of Byte Count Error	

NOTES:

In the unbuffered MSA, conditions causing the setting of bits 17 through 7 cause bit 16 of the normal status word to be set.

In the buffered MSA, conditions causing the setting of bits 22 through 18 cause bit 16 of the normal status word to be set.

Figure A–3. Auxiliary Status Word Format

A.4.1. Auxiliary Status Word Indicator

Bit position 35 is the auxiliary status word indicator and is always zero. Bit 35 indicates that the status word transmitted to the processor is the auxiliary status word.

A.4.2. Service In/Data In Received (SDIR)

The service in/data in received field (SDIR), in bit position 23, is used only in the buffered MSA status word. The SDIR field indicates that at least one SERVICE IN or DATA IN has been detected in the MSA.

A.4.3. Error Detection Code (EDC)

The error detection code field (EDC), in bit positions 22 through 18, is used in the buffered MSA only. The EDC field indicates that portion of the hardware which has detected an error condition. If any bit is set in this field, the WPC bit (bit 9) will also be set.

A.4.4. Control Unit Nonoperational (NOP)

The control unit nonoperational (NOP) bit, in bit position 17, is set when the disc unit SCU is offline to the MSA or when power to the disc unit or SCU is off.

A.4.5. Input Parity Error (IPE)

The input parity error bit (IPE), in bit position 16, is set in response to an address/status parity error or a data parity error. An address/status parity error is the detection of a parity error on the input bus during transmission of the device address byte or status byte from the device. A data parity error is the detection of a parity error on the input bus from the device to the MSA during an input data transfer sequence. Once the error is detected, the MSA rejects the data byte by responding with the COMMAND OUT signal to the SERVICE IN signal, thereby terminating data transfers for the remainder of the command in progress. The error byte is not transmitted to the processor and the position of the last valid data byte is indicated by bits 3 through 0 of the auxiliary status word.

A.4.6. Time Check (TCK)

The time check bit (TCK) in bit position 15, is set when either of two conditions occur on the device bus.

- After the start of an operation, if the device interface is inactive for at least 24.3 seconds (in the unbuffered MSA) or 50 seconds (in the buffered MSA), the time check bit is set. This inactivity can occur on a signal exchange when the device control unit or MSA fails to respond with the appropriate signal within the allotted time period. In addition, once the initial selection sequence is complete, the device control could fail to respond with a data or status request.

- Once the interface is activated, if it fails to go inactive within at least 70 microseconds, the time check bit is set. Each time the interface goes active, the delay is retriggered. This means a single signal exchange should not require more than the specified time period.

A.4.7. Control Line Error (CLE)

A control line error bit (CLE), in bit position 14, is set when two or more inbound control lines are activated at the device interface simultaneously. The input control lines which are checked include: SELECT IN, ADDRESS IN, STATUS IN, SERVICE IN, and DATA IN.

A.4.8. Address Compare Error (ACP)

The address compare error bit (ACP), in bit position 13, is set when the 8-bit device address transmitted to the device does not agree with that returned by the device during the initial selection sequence.

A.4.9. Translate Check (XCK)

The translate check bit (XCK), in bit position 12, is set if the translate mode is active (bit 19 in the function word is set to 1), if the current operation is an input forward or input backward, and if a data byte received by the MSA from the device is outside the range of the activated translator. If this bit is set, the E bit of the status word is also set.

A.4.10. Late Acknowledge (LA)

The late acknowledge bit (LA), in bit position 11, is set when a response from the processor to a data request is late (processor late).

A.4.11. Stall Check (SC)

The stall check bit (SC), in bit position 10, is set when a stall condition is detected within the MSA. The stall condition occurs when a connection is commenced but not completed at the processor interface, that is, if the device interface is not activated within about half a second after the MSA is activated as a result of receiving a forced external function. For example, if the processor external function buffer terminates with the command chaining flag set to 1 in the last function word transmitted, the MSA will expect another function word to follow and will not activate the device interface until all function words have been received. Therefore, if another processor buffer is not activated within the allotted time of approximately 0.5 second, the MSA terminates operation with the stall check indication.

A.4.12. Word Parity Check (WPC)

The word parity check bit (WPC), in bit position 9, indicates that a parity error was detected in a function or output data word received from the processor. Bit 9 is set to 0 if checking hardware is not installed (or is disabled by patch wiring).

A.4.13. Invalid Sequence (IS)

The invalid sequence bit (IS), in bit position 8, is set if any of the following conditions are encountered:

- When a function word which specifies the MSA Jump command also specifies chaining.

- When a Jump command is issued as a single function.

- When the Jump command address (bits 26 through 24) in a function word which specifies the MSA Jump command is greater than or equal to the address of the jump function word.

- When a jump loop does not contain functions requiring input or output transfers.

- When the chain flag was specified on the eight function (the second function if F1323–00, function buffer expansion, is not installed) received from the processor in a chain.

A.4.14. Invalid Command (IC)

The invalid command bit (IC), in bit position 7, is set:

- When a function word specifies a Test command in bits 15 through 8 when chaining is specified (bit 21 is 1). A Test command, then, cannot be specified in a chain except in the last function position.

- When a function word specifies translation mode, and the MSA does not have the appropriate feature.

- When codes issued to the MSA ($F1_{16}$) that are outside the command set, as specified by MSA functions, are detected.

- When the S flag (bit 23) and M flag (bit 20) bits are set to one in the same function word.

- When the write translate table command is issued and no feature existed.

A.4.15. Function Address (FA)

The function address field (FA), in bit positions 6 through 4, indicates the internal address of the function executed last in the MSA. The address relates to the function buffer located in the MSA which stores chained functions. The FA field is intended as an aid to error recovery when errors occur in the middle of a chain and, except as noted below, indicates the address of the last function executed. Execution of the MSA test function does not alter the stored count.

A.4.16. Magnitude of Byte Count Error (M)

The magnitude of the byte count error field (M), in bit positions 3 through 0, indicates the position of the last significant information byte in the last word transmitted to or received from the device. The ABC bit in the normal status word indicates that the last data word was not fully packed or unpacked. The M field is ignored unless the ABC bit is set in the normal status word.

A.5. DATA WORD FORMATS

A data word can be in one of three formats designated A, B, or C: A is the quarter-word format, B is the six-bit packed format, and C is the eight-bit packed format.

The word assembly and disassembly process is shown in Figure A–4 for each of the three formats. The relationship between the bit position of the 36-bit processor word and the bit positions of the eight-bit device bytes is also shown for the forward direction. For operation in the backward direction, the order of word transfer and order of byte transfer are reversed.

A.5.1. Quarter-Word, Format A

An output data transfer with format A can stop on any quarter-word boundary. During data transfer, bits 8, 17, 26, and 35 of the 36-bit processor data word are monitored in the MSA for a 1 bit, indicating stop control. The first stop bit terminates data transfer to the SCU. The byte in the same quarter-word as the stop control is not to be transmitted. If stop control is used for the termination of data transfers, the 0 bit in the current function word is ignored. On an input data transfer, these bits are zero.

A.5.2. Six-Bit Packed, Format B

When performing an output operation in format B, the MSA's word-to-byte converter forces bits 0 and 1 of each byte to zero (see Figure A–4b). If translation is not specified, bit 0 and 1 of each byte leave the MSA as zeros.

When performing an input operation in format B, bit 0 and 1 are ignored by the MSA's byte-to-word converter (see Figure A–4b).

A.5.3. Eight-Bit Packed, Format C

When operating in format C, output buffer data may end in the middle of a byte (byte 5, Figure A–4c). If the SCU is presenting a request for output data, the remainder of the byte is zerofilled (bits 4–7) by the MSA and transmitted and recorded.

The reverse situation occurs on an input data transfer when this record is read. The SCU sends a zerofilled byte 5 and presents end status to the MSA. The MSA completes assembly of word 1 (Figure A–4c) and discards the remaining bits (bits 4–7) of byte 5.

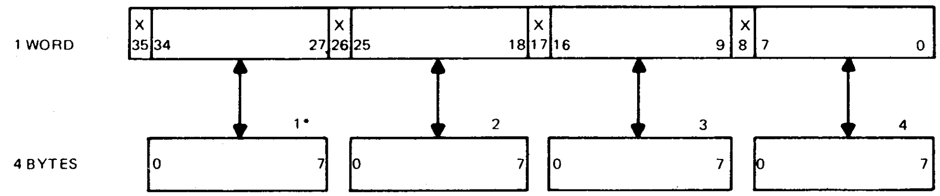

NOTE:

Bits 35, 26, 17, and 8 are used for stop control on output operations and forced to binary 0 on input operations.

Figure A–4a. Format A (Quarter Word)

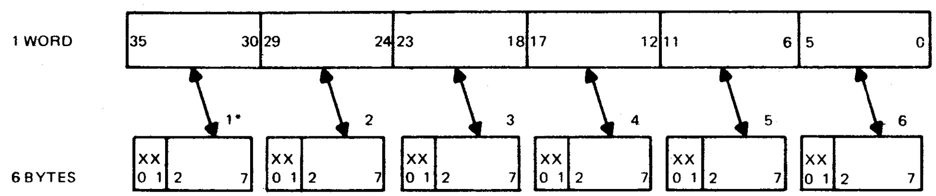

Figure A–4b. Format B (6-Bit Packed)

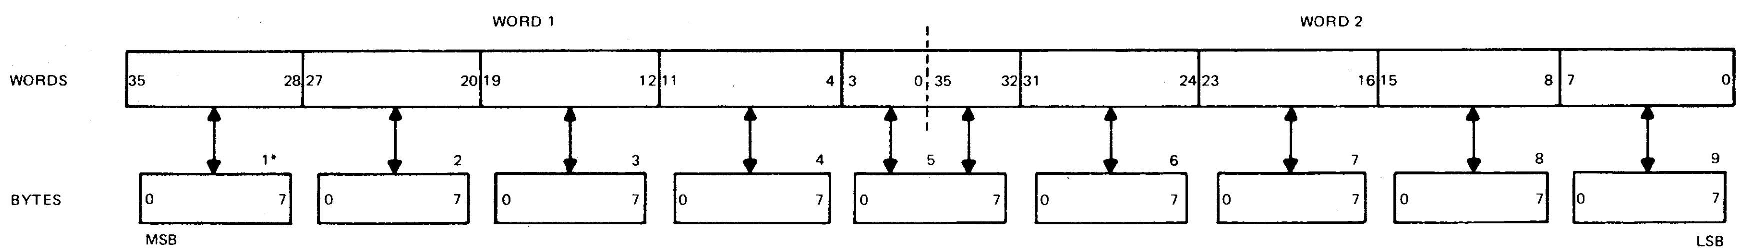

Figure A–4c. Format C (8-Bit Packed)

*Numbers on arrows indicate the order of byte transfer.

Figure A–4. Data Word Formats

USER COMMENT SHEET

Comments concerning this manual may be made in the space provided below. Please fill in the requested information.

System: ____________________

Manual Title: ____________________

UP No: __________ Revision No: __________ Update: __________

Name of User: ____________________

Address of User: ____________________

Comments:

FOLD

FIRST CLASS

PERMIT NO. 21

BLUE BELL, PA.

BUSINESS REPLY MAIL

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

POSTAGE WILL BE PAID BY

SPERRY UNIVAC

SYSTEMS SUPPORT
ATTN: INFORMATION SERVICES M.S. 4533
P.O. BOX 3942
ST. PAUL, MINNESOTA 55165

CUT

FOLD