SPERRY RAND

UNIVAC 1100 SERIES

8414/8411 DISC SUBSYSTEMS

PROGRAMMER REFERENCE

UP-7905

CONTENTS

FIGURES

TABLES

1. INTRODUCTION

This manual contains the information and procedures for the programmer when using the UNIVAC 8414 Disc Subsystem with a 36-bit, word-oriented processor. Information contained herein also applies to the UNIVAC 8411 Disc Subsystem since the two subsystems differ only in storage capacity, bit density of the discs, and transfer rates. These differences are treated separately in Appendix A for the 8411 subsystem.

This manual is divided into the following sections:

- Subsystem Description
- Programming
- Programming Example

The UNIVAC 8414 Disc Subsystem (Figure 1–1) provides 36-bit, word-oriented processors with millisecond random or sequential access to large data files for online use. Since the disc subsystem is an 8-bit, byte-oriented device, a UNIVAC 0961-02 or 0961-05 Multi-Subsystem Adapter (MSA) is required, to provide compatibility. See Figure 1–2. Information pertaining to the MSA is found in *UNIVAC 1100 Series Multi-Subsystem Adapter Programmer/Operator Reference, UP-7890* (current version). The UNIVAC 8414 Disc Subsystem, sometimes referred to as a disc file, consists of a disc drive, a disc pack, and a control unit.

The disc drive with a disc pack installed performs the following:

- Receives control signals from the control unit.
- Notifies the control unit of various conditions within a disc drive as to when it is selected and its operation.
- Accepts data to be written on a disc pack from the control unit.
- Sends data read from a disc pack to the control unit.

The control unit performs the following:

- Receives and interprets command information sent from the processor by way of the MSA.
- Accepts data from an addressed disc drive with its associated disc pack and makes that data available to the processor.
- Requests data from a processor and controls the writing of that data on the disc pack in the addressed disc drive.
- Checks for errors.

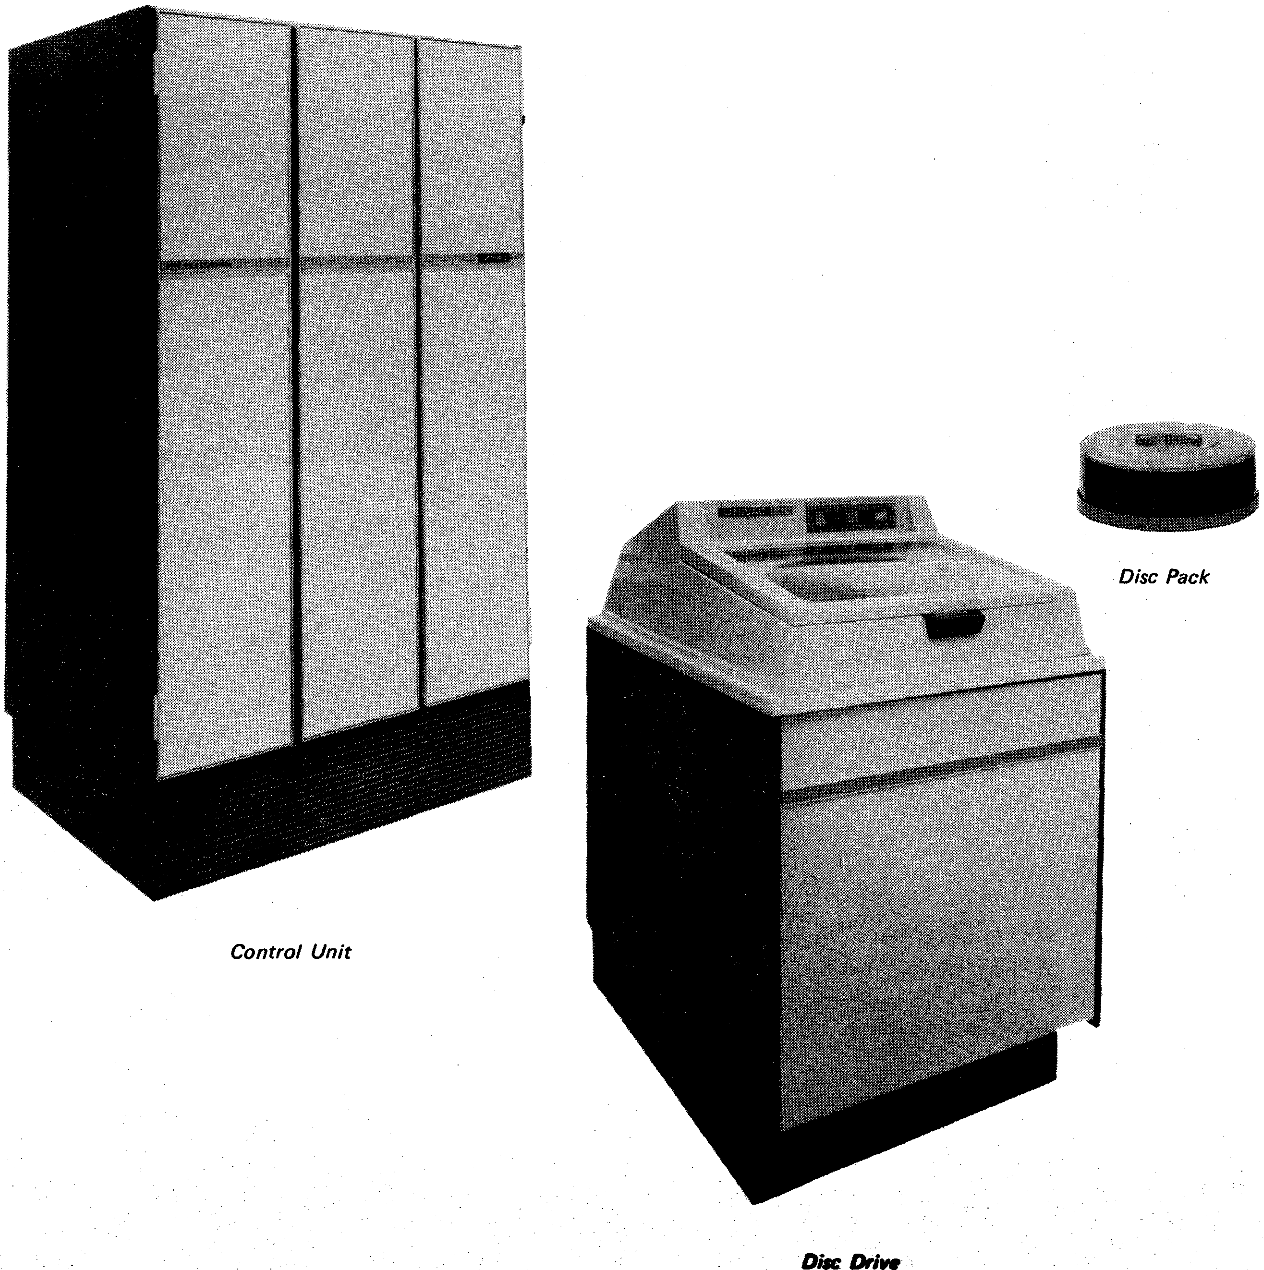

Figure 1–1. UNIVAC 8414 Disc Subsystem

The Multi-Subsystem Adapter performs the following:

- Enables 8-bit, byte-oriented disc subsystems to communicate with 36-bit, word-oriented processors.
- Communicates over normal I/O channels in internally specified index mode.
- Accepts function words from a processor and initiates commands by way of its initial selection sequence to the UNIVAC 8414 Disc Subsystem.

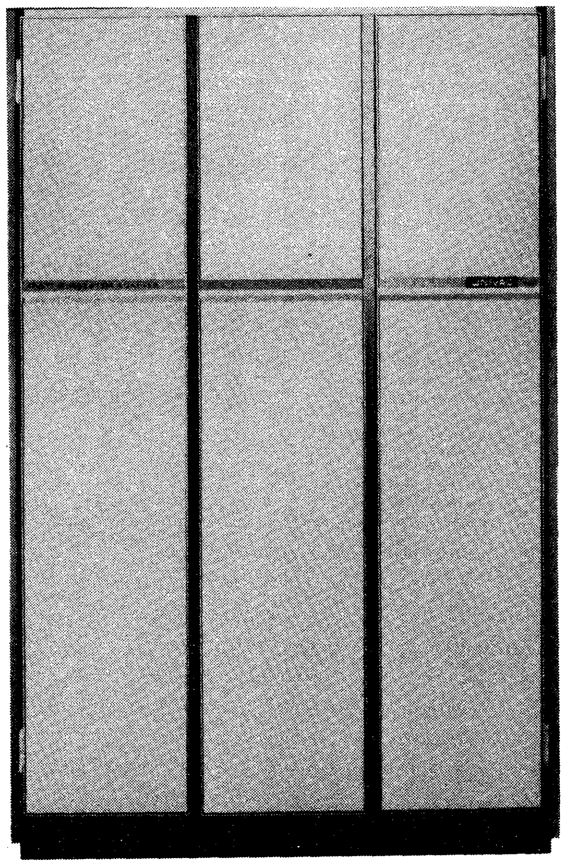

Figure 1–2. UNIVAC 0961 Multi-Subsystem Adapter

2. SUBSYSTEM DESCRIPTION

2.1. UNIVAC 8414 DISC SUBSYSTEM

The UNIVAC 8414 Disc Subsystem is a random access storage unit using a disc drive with removable and interchangeable disc packs. A control unit provides the instructions and control for reading and writing information on the surfaces of the discs.

A summary of the UNIVAC 8414 Disc Subsystem characteristics is given in Table 2–1.

DISC DRIVES PER SUBSYSTEM	2 to 8
DISC DRIVES PER CABINET	1
READ/WRITE HEAD ACCESSOR MECHANISMS	1
READ/WRITE HEADS PER DISC DRIVE	20
DISC PER PACK	11
RECORDING SURFACES PER DISC PACK	20
TRACKS PER DISC SURFACE	203
MAXIMUM EIGHT-BIT BYTES PER TRACK	7,294
DATA CAPACITY (EIGHT-BIT BYTES) PER DISC PACK (1)	29,176,000
TRACK DENSITY	100 per inch
BIT DENSITY	2200 bits per inch
DATA TRANSFER RATE	312,000 bytes per second
DISC SPEED	2400 rpm
DISC DIAMETER	14 inches
VERTICAL SPACING BETWEEN DISCS	0.350 inch
PARITY	Odd byte parity on command, address, status, sense, and data bytes. Cyclic parity on all record fields

(1) Data capacity, which is based on 200 tracks per disc surface, allows for the possibility of sixty unusable tracks per disc pack.

Table 2–1. Characteristics of UNIVAC 8414 Disc Subsystem

2.2. UNIVAC 0961 MULTI-SUBSYSTEM ADAPTER

The Multi-Subsystem Adapter (MSA) is connected to the processor through an I/O channel. It converts 8-bit byte information from a disc subsystem control unit into 36-bit words suitable for the processors. The MSA accepts and transfers four basic word configurations in its operation. They are the function word, the status word, the auxiliary status word and the data word. These words are transmitted from the processor, disassembled into 8-bit bytes in the MSA, and then sent to a UNIVAC 8414 Disc Subsystem. In the opposite direction, the 8-bit bytes of the disc subsystem are assembled into 36-bit words in the MSA and transferred to the processor. Function words which are accepted by the MSA from the processor cause commands to be initiated to the disc subsystem by way of the MSA's initial selection sequence. Status byte sequences from the disc subsystem are transmitted back to the processor in the form of status words and auxiliary status words, by way of the MSA. Figure 2–1 is a brief flow diagram for the transfer and conversion of basic word-byte information. Since the function word is the primary control word required from the processor, the detailed format including the definition for each bit of the word is given in 3.7, all other words required are shown but bit descriptions are given in greater detail in the *UNIVAC 1100 Series Multi-Subsystem Adapter Programmer Reference, UP-7890* (current version).

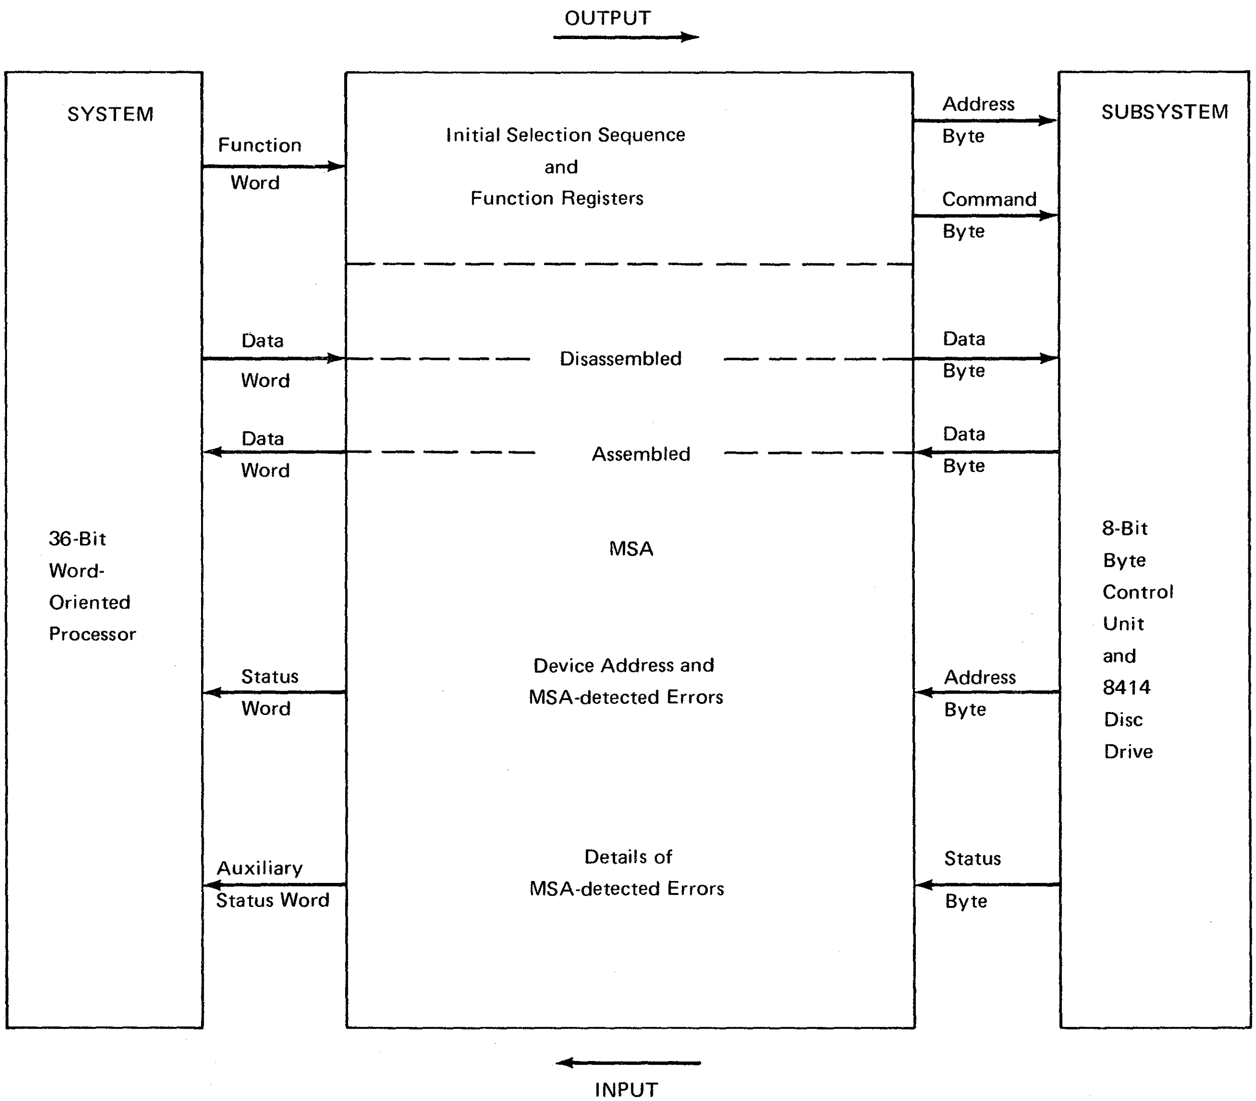

Figure 2–1. Basic Word and Byte Transfers

2.3. CONFIGURATIONS

The minimum operating configuration comprises a Multi-Subsystem Adapter, a control unit, and two disc drives. For single-channel operation, expansion to a maximum configuration of eight disc drives can be accomplished with one control unit. A dual channel operating subsystem would include using a Multi-Subsystem Adapter, a Multi-Subsystem Adapter expansion, a control unit containing a dual-channel feature and from two to eight disc drives. This configuration provides access to the control unit from two processor I/O channels. A dual access configuration consists of a Multi-Subsystem Adapter and expansion feature, two control units, and from two to eight disc drives per control unit. This configuration provides simultaneous read/read, read/write, or write/write operation.

Interchangeable, industry-compatible magnetic disc packs with a storage capacity of 29.17 million 8-bit bytes are available for use with the disc drives.

The minimum configuration (capacity of 58.34 million bytes) can be expanded by adding up to six additional disc drive units, for an online capacity of 233.36 million bytes. Special features are described in the following paragraphs along with Figure 2–2 which shows diagrams for single- and dual-channel operation and dual-access operation.

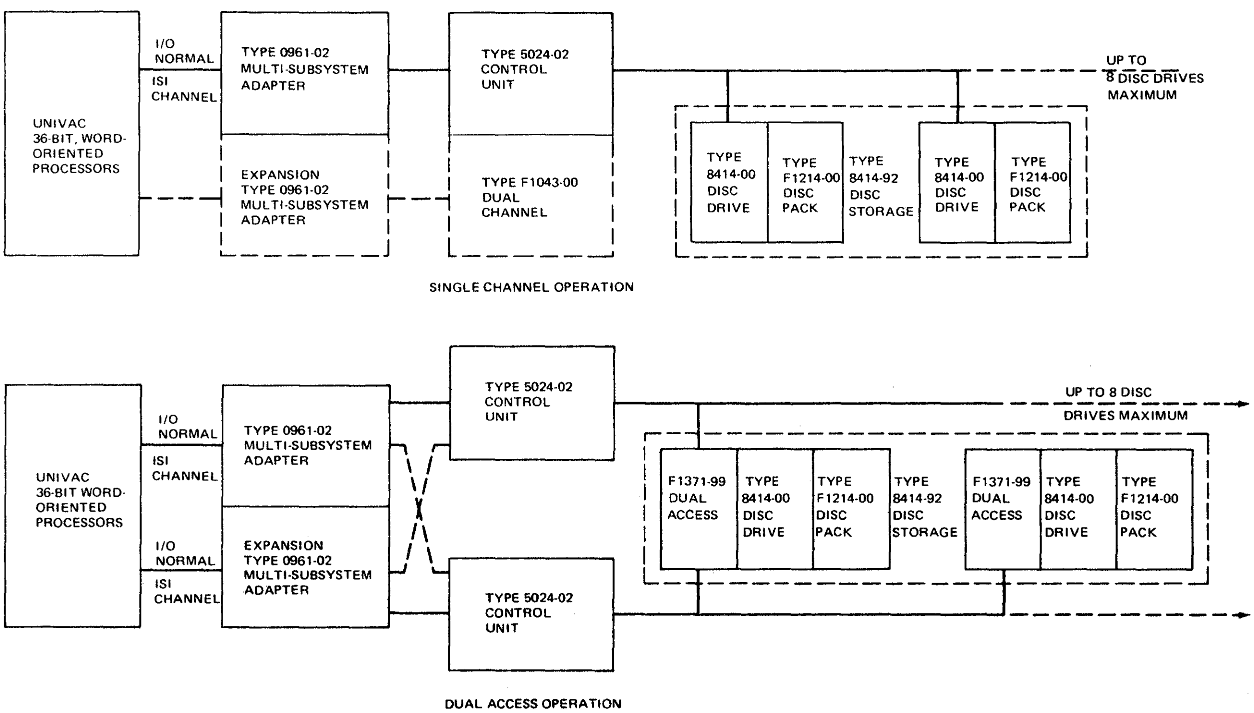

Figure 2–2. Operating Configurations for the UNIVAC 8414 Disc Subsystem

2.3.1. Basic Components And Optional Features

Table 2–2 lists the basic components and optional features of the UNIVAC 8414 Disc Subsystem.

COMPONENT/OPTIONAL FEATURE	TYPE NUMBER		REQUIREMENT PER SUBSYSTEM
	60 Hz	50 Hz	
Control Unit Dual Channel	5024-02 F1043-00	5024-05 F1043-00	1 1 optional
Disc Drive Unit Dual Access	8414-00 F1371-99	8414-01 F1371-99	2, 4, 6, or 8* 1 per disc drive unit, optional
Disc Pack	F1214-00	F1214-00	1 per disc drive unit, minimum

* Two Type 8414-00/01 Disc Drives are provided as a basic unit of disc storage and identified as Type 8414-92.
Four Type 8414-00/01 Disc Drives are provided as a basic unit of disc storage and identified as Type 8414-94.
Six Type 8414-00/01 Disc Drives are provided as a basic unit of disc storage and identified as Type 8414-96.
Eight Type 8414-00/01 Disc Drives are provided as a basic unit of disc storage and identified as Type 8414-98.
One Type 8414-85 disc drive is provided for single unit configuration expansion.

Table 2–2. Basic Components and Features of UNIVAC 8414 Disc Subsystem

2.3.2. Disc Storage Optional Features

Table 2–3 lists the type numbers required to expand the particular disc drive basic units available to the disc subsystem.

OPTIONAL FEATURE AND TYPE NUMBER	PURPOSE
Disc Drive 8414-91	Expands basic disc storage Type 8414-92 by one or two disc drive units.
Disc Drive 8414-90	Expands basic disc storage Type 8414-94 by one or two disc drive units.
Disc Drive 8414-89	Expands basic disc storage Type 8414-96 by one or two disc drive units.

Table 2–3. Disc Storage Optional Features

2.4. DISC SUBSYSTEM COMPONENTS

A disc subsystem comprises from two to eight disc drives, interchangeable industry-compatible disc packs, and a control unit capable of handling from two to eight disc drives. The disc drive provides the mechanical drive for the disc pack; the disc pack, consisting chiefly of magnetic recording discs, is the storage element, and the control unit provides the electronic direction and control for writing and retrieving data.

2.4.1. Disc Drive Unit

When instructed to do so through a control unit, the disc drive unit performs read and write operations on the accessed disc surface of a disc pack. The disc drive contains the motors and drive mechanisms required to control the rotation of the discs at 40 revolutions a second. The drive unit also contains the support mechanism required for the read/write recording heads.

2.4.2. Disc Pack

The disc pack is the storage medium of the disc drive unit. It consists primarily of eleven disc platters mounted on a vertical shaft. Data is written on the inside surfaces of each disc. No data is recorded on the upper side of the top disc or the under side of the bottom disc. There are circular protective plates above the top disc and below the bottom disc to protect the disc pack from mechanical injury.

2.4.2.1. ACCESS MECHANISM

The access mechanism consists of an access carriage on which two columns of read/write head arm assemblies are mounted. Each column contains ten assemblies. The read/write heads are moved by a hydraulic accessor system horizontally between the tracks, beginning near the periphery and ending at the center hub of each disc. The heads are moved directly from one track to another without returning to a home position.

2.4.2.2. FILE ORGANIZATION AND ADDRESSING FORMAT

Access to the data on the disc pack is obtained by positioning the comb-like array of read/write heads to any one of the 203 disc positions. Since all read/write heads are in the same vertical plane, corresponding tracks on the several recording surfaces are accessible without movement of the access mechanism. These tracks are numbered 0 to 19, from top to bottom. These vertical tracks can be regarded as a cylinder of data; and switching between the tracks of the same cylinder is accomplished electronically. Figure 2–3 is a cutaway view showing the entire disc pack. Each disc consists of 203 concentric cylinders numbered 000 to 202 starting from the outside. The address of an individual track in a given disc pack therefore consists of a cylinder number and a read/write head (track) number. The addressing format is also shown in Figure 2–3.

2.4.2.3. STORAGE CAPACITY

The byte capacity is based on 200 cylinders per disc pack. The three remaining cylinders are spares provided to ensure that the stated capacity is maintained should some tracks or cylinders become permanently damaged. Because each record has some nondata characters, such as its address, the net data storage capacity of the tracks may vary. The major loss of storage area is due to the gaps between records.

Based on one record per track, not including home address and first record (R_0), the capacity of each track is a function of the number of equal length records on that track.

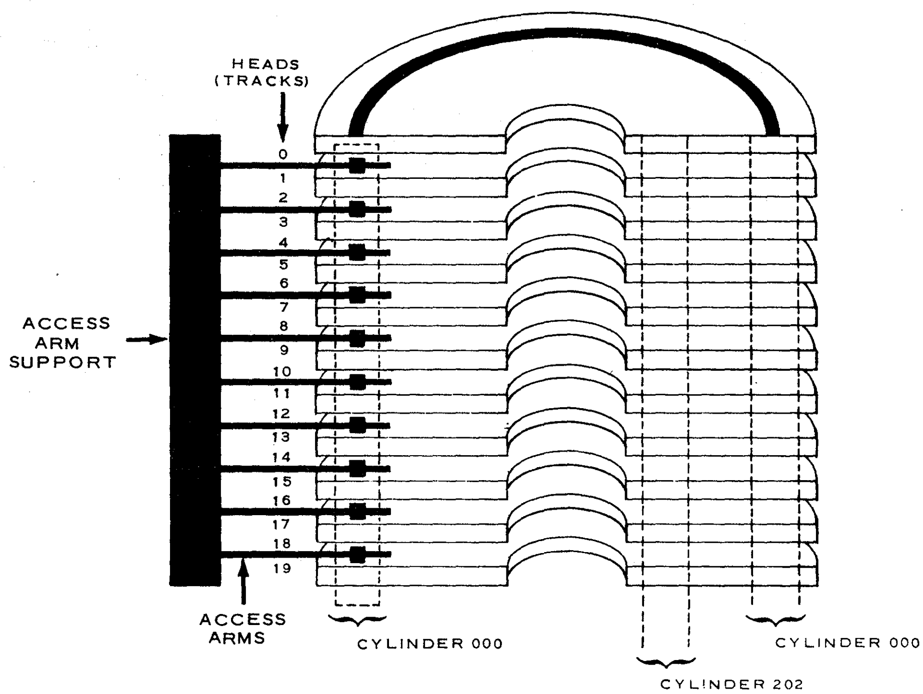

	U	C1	C2	H1	H2
BIT	01234567	01234567	01234567	01234567	01234567
	BYTE 1	BYTE 2	BYTE 3	BYTE 4	BYTE 5

Legend:

U = Disc Drive Unit Address (Bits 0–4, control unit; bits 5–7, disc drive unit)

C1 and C2 = Cylinder Address; C1 must always be 0; C2 is the cylinder identification (from 000 to 202_{10})

H1 and H2 = Head Address; H1 must always be 0; H2 is the track digit number, $0–19_{10}$. Bit position 7 of byte 5 is the least significant bit.

Figure 2–3. File Organization and Addressing Format

2.4.2.4. DISC TRACK AND RECORD FORMATS

The contents of each track consists of a home address, a track descriptor record (R_0), and one or more data records. A table relating the number of bytes per record to the number of records per track is given in Table 2–4. Each data record has an identifying address marker. The disc track and record formats are shown in Figure 2–4.

Maximum Number of Data Bytes per Record	NUMBER OF EQUAL LENGTH RECORDS PER TRACK																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Without Key Field	7294	3520	2298	1693	1332	1092	921	793	694	615	550	496	450	411	377	347	321	298	276	258
With Key Field	7249	3476	2254	1649	1288	1049	877	750	650	571	506	452	407	368	333	304	277	254	233	215

Maximum Number of Data Bytes per Record	NUMBER OF EQUAL LENGTH RECORDS PER TRACK									
	21	22	23	24	25	26	27	28	29	30
Without Key Field	241	226	211	199	187	176	166	157	148	139
With Key Field	198	183	168	156	144	133	123	114	105	96

NOTE: All bytes are read left to right.

Table 2–4. Bytes Per Record, Function of Records Per Track

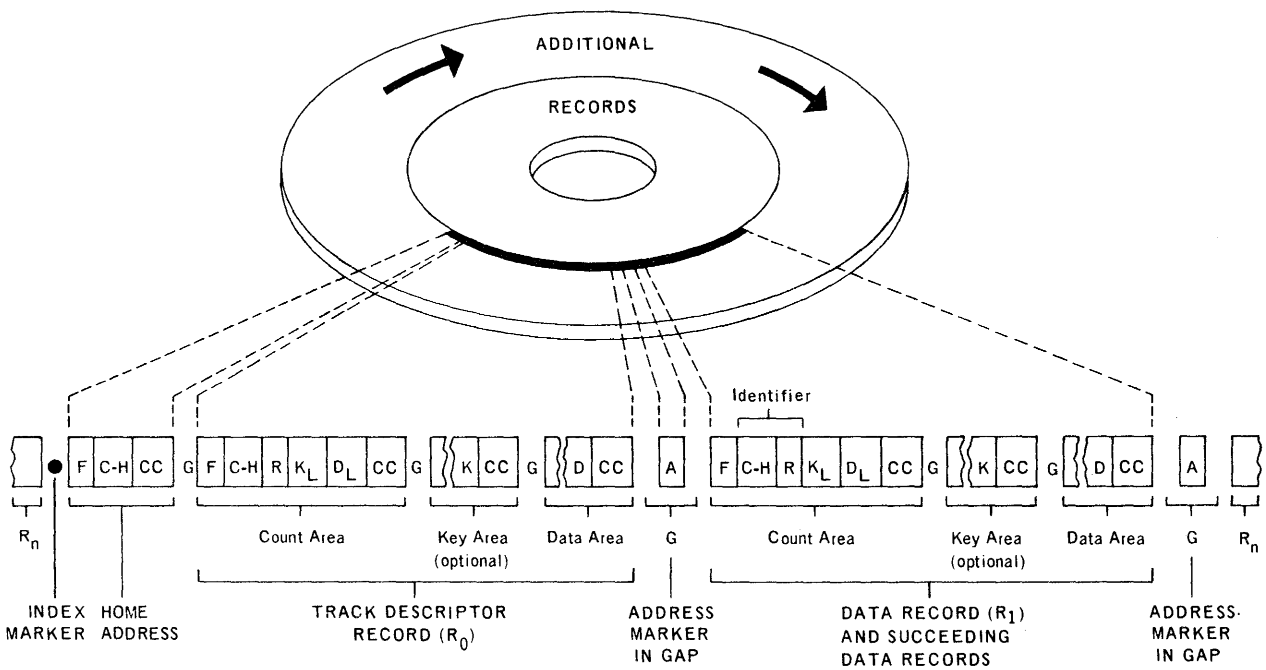

Legend:

F = FLAG (ONE BYTE)
C–H = CYLINDER AND HEAD (FOUR BYTES)
CC = CYCLIC CHECK (TWO BYTES)
G = GAP (PROPORTIONATE IN LENGTH TO PRECEDING FIELD)
R = RECORD NUMBER (ONE BYTE)
K_L = KEY LENGTH (ONE BYTE)
D_L = DATA LENGTH (TWO BYTES)
K = KEY
D = DATA
A = ADDRESS MARKER (TWO BYTES)
R_n = SUCCEEDING DATA RECORDS

Figure 2–4. Disc Track and Record Formats

2.4.2.4.1. Index Marker

An index pulse is generated once every disc revolution using a notch at the bottom of the disc pack assembly. As all tracks are aligned to have the same beginning point, this marker indicates the physical beginning of every track on the disc pack.

2.4.2.4.2. Home Address

The home address consists of a flag byte (also written for record R_0), address bytes, and check bytes. The status and physical location of a track is defined in the home address. The seven bytes of the home address are:

- The flag byte (F), which has the following significance:

BIT	FUNCTION
0–5	0's
6	0 denotes operative track.
	1 denotes defective track.
7	0 denotes primary track.
	1 denotes alternate track.

- The four cylinder and head address bytes (C–H), which identify one of the 203 cylinders and one of the read/write heads.

- The two cyclic check bytes (CC), which are automatically computed and appended to the record by the control unit whenever a home address is written. They are recomputed and compared to the originally written bytes to check the validity of the address whenever it is read.

2.4.2.4.3. Record

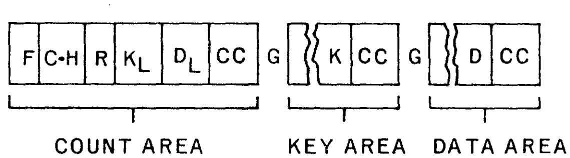

The number and the length of data records are limited only by the track capacity; however, all elements within each record must conform to the format shown in Figure 2–4.

The first record (R_0) after the home address is designed to facilitate the use of an alternate track address if the addressed track cannot be used. This record is not preceded by an address marker and usually does not contain a key area. The flag byte of record R_0 is always the same as the home address flag byte described in 2.4.2.4.2.

The address marker (A) consists of two bytes which identify the beginning of each record except record R_0.

- Count Area

The count area is used to describe the physical location of the record and the length of the key and data areas. The bytes in the count area are:

(1) The flag byte (F), supplied by the control unit as records R_1, R_2,.....R_n are written. The flag bytes for record R_0 is the same as the home address flag byte. The flag byte has the following significance:

BIT	FUNCTION
0	0 indicates even-numbered records.
	1 indicates odd-numbered records.
1	0 indicates normal (nonoverflow) records.
	1 indicates record overflows to the next track. This flag bit is set as special-count-key-and-data commands write the count field of each portion of the overflow record except the last.
2–5	0's
6	0 denotes operative track.
	1 denotes defective track.
7	0 denotes primary track.
	1 denotes alternate track.

(2) The address bytes (C–H), identify one of the 203 cylinders and one of the 20 read/write heads. If the track is bad, the cylinder and head numbers of an alternative track are specified in the count area.

(3) The record number byte (R), is the sequential number of the record on the track. Since the track descriptor record is R_0, its record number byte contains all binary 0's.

NOTE: The cylinder number, head number (C–H), and record number (R) located in the count area are designated the identifier; see Figure 2–4.

(4) The key length byte (K_L) indicates the number of bytes, exclusive of the check bytes, in the key area of the record. Lengths from 0 (no key area) to 255 bytes can be indicated by the key length byte. Since record R_0 usually does not contain a key area, the key length byte usually is all binary 0's and the count area is followed by the data area.

(5) The data length bytes (D_L) indicate the number of bytes, exclusive of check bytes, in the data area of the record. If the data length bytes are all binary 0's, the data area contains one byte of all binary 0's and the two check bytes. When a no-data record is read, an end of file is sent to the channel from the control unit. Lengths from 0 (no data to be transferred) to the maximum length of 7294 bytes can be indicated by the data length bytes.

(6) The cyclic check bytes (CC) are computed and written when a count area is written. These are read back and compared to recomputed check bytes whenever a count area is read. Thus error checking is performed for the count area read.

- Key Area

  Used as an identifier for a specific data area on a track. The key area is used to permit, during a single disc revolution, searching for a record and reading or writing it. Two cyclic check bytes allow error checking of the information read from the key area. This area is usually not used in record R_0.

- Data Area

  Contains information written to be available on demand. Two cyclic check bytes allow error checking of the information read from the data area.

2.4.3. Control Unit

The control unit, which is an integral part of a subsystem is connected to the MSA. The control unit can control up to eight disc drives and can provide a one-byte wide data path to the MSA and a one-bit wide data path to the disc units. Data can be transferred over these data paths in either direction on a nonconcurrent basis.

The general functions of a control unit have already been listed. In particular a control unit performs the following:

- Assumes and relinquishes control of the interface between MSA and disc subsystem.

- Checks incoming data for odd parity and generates the parity bit for all outgoing bytes.

- Initiates an operation at an addressed disc unit.

- Assures correct data and address recording during read commands by a combination of a parity bit and check bytes for error checking.

3. PROGRAMMING

3.1. GENERAL

This section is intended to familiarize the programmer with the following:

- 8414 Disc Subsystem
  - Formats of executable commands, status bytes, and sense data bytes
  - Programming considerations
  - Timing factors of the subsystem
- Multi-Subsystem Adapter
  - Word Formats
  - MSA Functions

3.2. COMMAND BYTES

The UNIVAC 8414 Disc Subsystem recognizes the command bytes listed in Table 3–1 as valid commands. Each command byte is checked for parity and validity before being accepted by the control unit. Status bytes generated in response to command bytes are described in 3.3. Disc operation is controlled by the function word from the processor. This word is described in 3.7.1.

3.2.1. Command Chaining

The function word contains a command chain bit (bit 21, C of Multi-Subsystem Adapter function word) used for command chaining. When the command chain bit is a 1, the operation specified by the command byte in the immediately following function word is initiated. After receipt of device end and channel end signals by the processor I/O channel for normal completion of the current operation, the next operation is initiated. When the command chain bit is not a 1, the chain is terminated. Chaining control takes place in the MSA with eight or fewer control functions.

COMMAND	CODE							
SEEK	0	1	2	3	4	5	6	7 ①
	15	14	13	12	11	10	9	8 ②
Seek	0	0	0	0	0	1	1	1
Seek-head	0	0	0	1	1	0	1	1
Seek-cylinder	0	0	0	0	1	0	1	1
WRITE								
Write-home-address	0	0	0	1	1	0	0	1
Write-TD-record	0	0	0	1	0	1	0	1
Write-count-key-and-data	0	0	0	1	1	1	0	1
Write-special-count-key-and-data	0	0	0	0	0	0	0	1
Write-data	0	0	0	0	0	1	0	1
Write-key-and-data	0	0	0	0	1	1	0	1
READ	③							
Read-home-address	M	0	0	1	1	0	1	0
Read-TD-record	M	0	0	1	0	1	1	0
Read-count	M	0	0	1	0	0	1	0
Read-data	M	0	0	0	0	1	1	0
Read-key-and-data	M	0	0	0	1	1	1	0
Read-count-key-and-data	M	0	0	1	1	1	1	0
Initial-program-load	M	0	0	0	0	0	1	0

Table 3–1. Command Bytes (Part 1 of 3)

COMMAND	CODE							
SEARCH								
Search-home-address-equal	M	0	1	1	1	0	0	1
Search-ID-equal	M	0	1	1	0	0	0	1
Search-ID-high	M	1	0	1	0	0	0	1
Search-ID-equal-or-high	M	1	1	1	0	0	0	1
Search-key-equal	M	0	1	0	1	0	0	1
Search-key-high	M	1	0	0	1	0	0	1
Search-key-equal-or-high	M	1	1	0	1	0	0	1
Search-key-and-data-equal	M	0	1	0	1	1	0	1
Search-key-and-data-high	M	1	0	0	1	1	0	1
Search-key-and-data-equal-or-high	M	1	1	0	1	1	0	1
Continue-scan (see 3.2.5.9)								
SENSE								
Sense-I/O	0	0	0	0	0	1	0	0
Sense-reserve	1	1	1	1	0	1	0	0
Sense-release	1	1	0	1	0	1	0	0
MISCELLANEOUS								
Set-file-mask	0	0	0	1	1	1	1	1
Recalibrate	0	0	0	1	0	0	1	1
No-operation	0	0	0	0	0	0	1	1
Space-count	0	0	0	0	1	1	1	1
Erase	0	0	0	1	0	0	0	1
Test-I/O	0	0	0	0	0	0	0	0

Table 3–1. Command Bytes (Part 2 of 3)

① Bit positions in a byte, position 7 being the least significant bit position.

② Function word bit positions relative to a 36-bit word.

③ The M bit, when 0, establishes normal operation mode. The M bit, when 1, establishes multiple-track mode. This bit is ignored by the control unit on an initial-program-load command. When the M bit is set to 1 in the command, the disc unit upon encountering the index mark, increments the head register to switch to the next head. This M bit when set to 1 in a search-truncated command, (i.e., search-ID-equal command with P field set to 0, followed by a jump to the read/write command), and the track descriptor record (TDR) is used as a data record, enables the program to cascade down the cylinder switching to the next head after reaching the index mark. If the track descriptor record is not used as a data record, and the data length is 0 along with external interrupt status containing unit exception, disc transfer terminates. If the TDR does not have a data length of 0, the data within the TDR will then be presented, and the read/write, search-truncated, jump, and chain continue.

The M bit is bit 15 of the 36-bit function word or bit 0 of a byte to the control unit. This bit has no relationship to the M bit (bit 20) of the function word, which enables the MSA to execute a string of functions to different disc units without storing functions.

Table 3–1. Command Bytes (Part 3 of 3)

3.2.2. Seek Commands

The three seek commands (seek, seek-cylinder, and seek-head) initiate the transfer of six bytes of addressing information to the control unit as follows:

BYTE	SIGNIFICANCE	RANGE
0	Not used	Must be all 0's
1	Not used	Must be all 0's
2	Not used	Must be all 0's
3	Cylinder address	$0–202_{10}$
4	Not used	Must be all 0's
5	Head address	$0–19_{10}$

Seek command codes

COMMAND	BINARY CODE
Seek	0000 0111
Seek-cylinder	0000 1011
Seek-head	0001 1011

The seek and seek-cylinder commands cause the read/write head specified in byte 5 to be selected and cause the heads to be positioned at the cylinder specified in byte 3 of the address information. For these two commands, the channel end status is presented to the processor after the transfer of the six-byte address. Device end status is presented upon completion of accessor motion control. If no motion is required, or the function has the chain bit set; device end and channel end are presented together.

The seek-head command is used to select one of the 20 read/write heads, as specified in byte 5. To prevent accessor motion control, byte 3 is ignored by the disc subsystem. For this command, both channel end and device end statuses are presented to the processor after the transfer of the six-byte address.

While a command is being performed on one disc unit, another command chain may be issued to a second disc unit; such overlapping of commands is permitted.

3.2.3. Write Commands

The write commands cause information to be transferred from main storage, by way of the MSA, to a disc for storage as specified by the address byte. The control unit generates and appends the appropriate check bytes to the information.

After each write command is completed, channel end and device end are presented to the processor (device end causes chaining), subsequent commands in the chain other than write commands are received and executed by the disc subsystem. If a control unit is addressed before reaching the index mark following channel end and device end for a previous format-write command, a control unit busy, status modifier, and external status is presented to the processor. This action occurs because the control unit continues to erase the track until the index marker is reached, even though completion status for the format write is presented to the processor. When the index marker is reached, the operation is completed and a control unit end status is presented to the MSA and processor if control unit busy status was detected.

3.2.3.1. WRITE-HOME-ADDRESS

This command initiates the transfer of the five-byte home address from main storage and causes the address to be written on a track. This command must be preceded by an appropriate set-file-mask command (see 3.2.7.1).

3.2.3.2. WRITE-TRACK-DESCRIPTOR-RECORD

This command initiates the transfer of record R_0 from main storage and causes the record to be written on a track. The flag byte is the same as the home address flag byte and is provided by the control unit along with all check bytes. If key or data area bytes transferred from the channel are fewer than specified in their respective length bytes (K_L,D_L), the control unit causes these areas to be filled with 0's, thereby conforming the key and data areas to the lengths specified in the count area. The track descriptor record may be used as a data record.

This command must be preceded by an appropriate set-file-mask command (see 3.2.7.1) and chained from a write-home-address command or a search-home-address-equal command.

3.2.3.3. WRITE-COUNT-KEY-AND-DATA

This command initiates the transfer of a data record (R_1 or R_2 or R_3 or ... R_n) from main storage and causes the data record to be written on a track. The control unit hardware provides the address marker bytes, the flag byte, and the check bytes. If the key or data area bytes transferred from the MSA are fewer than specified in their respective length fields, the control unit causes these areas to be filled with 0's after the writing of all data transferred, thereby conforming the key and data areas to the length specified in the count area.

This command must be chained from a write-TD command, another write-count-key-and-data command, a search-ID-equal command, or a search-key-equal command.

3.2.3.4. WRITE-SPECIAL-COUNT-KEY-AND-DATA

This command allows the formatting of all segments (except the last) of a record which overflows to another track. An overflow record is one which continues beyond one track. A segment is that portion of an overflow record which is written on one track. This command differs from a write- count-key-and-data command only in that it causes a 1 to be written in bit 1 of the count area flag byte. This command must be issued for all segments but the last. The last segment is written with the issuance of a write-count-key-and-data command in which bit 1 of the flag byte is set to 0.

The written count areas pertain to the segment of the record on that track.

3.2.3.5. WRITE-DATA

This command initiates the transfer of data from main storage and causes the data to be written in the data area of a record. The data length is specified by the data length bytes in the count area of the record in the MSA. If the processor does not transfer a sufficient length of data, the control unit provides the 0's necessary to achieve conformity between data and data length bytes from the MSA.

When the data transferred is a segment of an overflow record, the transfer is not completed until a segment is written in which a 0 is in bit position 1 of the count area flag byte is written. After the first segment is transferred, the control unit waits for the index marker. When the index marker is reached, the control unit selects the next sequential head, checks the home address, and searches for the first address marker on the track. This operation continues until a segment containing a 0 in bit position 1 of its flag byte is detected. At the end of data transfers performed during this segment, the operation is terminated. Transfers do not have to start at the first segment of an overflow record. Any segment may be entered, but after the last segment is processed, no further head switching takes place. If the access control word is not large enough to supply data to all segments, zero's will be supplied to fill the last segment.

This command must be chained from a search-ID-equal command or a search-key-equal command. Note 2 of Figure 3–1 describes how the M bit couples with truncated search for continuous write within a cylinder.

3.2.3.6. WRITE-KEY-AND-DATA

This command initiates the transfer of a key and data bytes from main storage and causes them to be written in the key and data areas of a record. The lengths of the key and data transferred are specified by the key and data length bytes in the count area of the record. If the processor does not transfer sufficient lengths of key and data, the control unit provides the necessary 0's after writing all data transferred, with the result that conformity is achieved between the key and data transferred and their respective length fields in the count area.

When the data transferred is a segment of an overflow record, the transfer is not completed until a segment containing a 0 in bit position 1 of the count area flag byte is written. After the first segment is transferred, the control unit waits for the index marker. When the marker is reached, the control unit selects the next head in sequence, checks the home address, and searches for the first address marker on the track. This operation continues until a segment containing a 0 in bit position 1 of its flag byte is reached. At the end of data transfers performed during this segment, the operation is terminated. Transfers do not have to start at the first segment of any overflow record; any segment may be entered. After the last segment is processed, no further head switching takes place.

This command must be chained from a search-ID-equal command.

3.2.4. Read Commands

The read commands cause information to be transferred from disc storage as specified by the address byte, by way of the processor, to main storage. On all read commands, the control unit checks for errors, eliminates the check bytes, and presents the information to the processor by way of the MSA.

The M bit (bit 0) of the command byte allows operation within cylinder boundaries rather than head boundaries. If the M bit is 0, normal operation is indicated and head switching does not take place. If, however, the M bit is 1, the next read/write head in sequence is selected each time the index marker is reached. If the M bit is 1 and an index marker is detected during execution of the following read command, reading continues from the next head within the cylinder. Head switching continues until bit 2 (cylinder end) of sense data byte 1 is set to 1 and sent to the processor. A separate read command must be issued for each record. All read commands must be preceded by an appropriate search command. See note 3 of Figure 3–1.

Normal termination of a read command is indicated by channel end and device end bits set to 1 in the status byte. If an error occurs during a read, then the unit check bit of the status byte is 1 and is presented to the processor with channel end and device end statuses.

3.2.4.1. READ-HOME-ADDRESS

This command initiates the transfer of the five-byte home address to main storage.

3.2.4.2. READ-TRACK-DESCRIPTOR-RECORD

This command initiates the transfer of all three fields (count, without flag byte; key; and data) of the record R_0 to main storage.

3.2.4.3. READ-COUNT

This command initiates the transfer (to main storage) of eight bytes (the flag byte is omitted) of the count area of a data record following the next address marker read. Record R_0 count area cannot be read by this command, since record R_0 is not preceded by an address marker.

3.2.4.4. READ-DATA

This command initiates the transfer of the data area of a record to main storage.

When the data transferred is a segment of an overflow record, the transfer is not completed until a segment containing a 0 in bit position 1 of the count area flag byte is transferred. After the first segment is transferred (or the count and key areas are transferred from the first segment), the control unit waits for the index marker. When the index marker is reached, the control unit selects the next sequential head, checks the home address, and searches for the first address marker on the track. This operation continues until a segment containing a 0 in bit position 1 of its flag byte is detected.

At the end of data transfers performed during this segment, the operation is terminated. Transfers do not have to start at the first segment of an overflow record; any segment may be entered. After the last segment is read, no further head switching takes place.

3.2.4.5. READ-KEY-AND-DATA

This command initiates the transfer of the key and data areas of a record to main storage. This command can also be used for records in which the key area is nonexistent, in which case it produces the same result as a read-data command.

When the data transferred is a segment of an overflow record, the transfer is not completed until a segment containing a 0 in bit position 1 of the count area flag byte is transferred. After the first segment is transferred (only the key and data areas are transferred from the first segment), the control unit waits for the index marker. When the index marker is reached, the control unit selects the next sequential head, checks the home address, and searches for the first address marker on the track. This operation continues until a segment containing a 0 in bit position 1 of its flag byte is detected.

At the end of data transfers involving this segment, the operation is terminated. Transfers do not have to start at the first segment of an overflow record; any segment may be entered. After the last segment is read, no further head switching takes place.

3.2.4.6. READ-COUNT-KEY-AND-DATA

This command initiates the transfer to main storage of an entire data record, except for the flag byte, which is omitted from the count area, following the next address marker read. Record R_0 cannot be read by this command, since record R_0 is not preceded by an address marker.

When the data transferred is a segment of an overflow record, the transfer is not completed until a segment containing a 0 in bit position 1 of the count area flag byte is transferred. After the first segment is transferred (the count, key and data areas are transferred with the first segment), the control unit waits for the index marker. When the index marker is reached, the control unit selects the next head in sequence, checks the home address, and searches for the first address marker on the track. This operation continues until a segment containing a 0 in bit position 1 of its flag byte is detected. At the end of the data transfers performed during this segment, the operation is terminated. Transfers do not have to start at the first segment of an overflow record; any segment may be entered. After the last segment is read, no further head switching takes place.

3.2.4.7. INITIAL-PROGRAM-LOAD

This command initiates the transfer of information concerning initial loading of a program from a disc unit selected under processor and MSA control. This command is initiated by the setting of switches on both the processor and MSA operators' panels. Before issuing this command, the entire system must be reset.

When the subsystem receives the initial-program-load command, the following occur:

- Head 0 is selected, and the accessor is positioned at cylinder 000 on the selected disc unit.
- The data area of the first record following record R_0 is read and transferred to main storage.
- At the end of record R_1 data area, channel end and device end statuses are sent to the processor to terminate the operation.

3.2.5. Search Commands

The search command results in the transfer of information from the processor and MSA to the control unit as well as from the disc unit specified by the address byte to the control unit, where the two sets of information are compared byte by byte. The sense of the comparison is taken so that "is greater than" means that the value read from the disc is greater than that from the MSA. The information processed by the control unit is checked for errors as the search is executed. The M bit (bit 0) of the command byte is used to permit operation within cylinder boundaries rather than within head boundaries. If the M bit is 0, normal operation is indicated and head switching does not take place. If, however, the M bit is 1 and the index marker is detected during a subsequent search command, the next read/write head in sequence is selected. If the M bit is 1 and a search command is repeated, the search continues from the next head within the cylinder. Head switching continues until bit 2 (cylinder end) of sense byte 1 is set to 1 and sent to the processor by way of the MSA. A separate search command must be issued for each record.

The M bit should be set to 1 with discretion. For example, if M is 1 and the index marker is detected before an entire track has been searched, the first portion of the first track is ignored on the search. On continuous read/write chain, a search-to-equal should be accomplished using M bit set to 0 to establish the start of the data area selected.

3.2.5.1. SEARCH-HOME-ADDRESS-EQUAL

This command initiates the transfer of four bytes of home address information from both the processor (and MSA) and a disc unit to the control unit, where these bytes are compared for equality. The comparison is made byte by byte, and if a find is made, the status modifier bit is set to 1. The flag byte is not compared during this command.

3.2.5.2. SEARCH-ID-EQUAL

This command initiates the transfer of the cylinder number, head number, and record number bytes from both the processor (and MSA) and a disc unit to the control unit, where these bytes are compared for equality. This command is used to permit the locating of a record. The identifier (ID) to be compared is the one following the next address marker or index marker. This command is repeated until a find is made or until two index markers have been detected. This command with no ID supplied (P field set to 0) and the control unit feature enables simulated continuous read/write chains.

3.2.5.3. SEARCH-ID-HIGH

This command is similar to a search-ID-equal command except that a find is made when the ID on the disc is greater than the ID presented by the processor by way of the MSA.

3.2.5.4. SEARCH-ID-EQUAL-OR-HIGH

This command is similar to a search-ID-equal command except that a find is made when the identifier on the disc is equal to or greater than the identifier presented by the processor and MSA.

3.2.5.5. SEARCH-KEY-EQUAL

This command initiates the transfer of the key area bytes from both the MSA and a disc unit to the control unit, where they are compared for equality. This command is used to permit the searching and reading or writing of the data area on a single disc revolution. If this command is chained from a search-ID command that had a find on record R_0, then the key area to be compared is either the one following the next address marker or the key area of record R_0. This command and the other search-key commands do not return a status modifier bit to the processor if the key length of the record examined is 0.

3.2.5.6. SEARCH-KEY-HIGH

This command is similar to a search-key-equal command except that a find is made when the key area on the disc is greater than the key area presented by the processor (and MSA).

3.2.5.7. SEARCH-KEY-EQUAL-OR-HIGH

This command is similar to a search-key-equal command except that a find is made when the key area on the disc is equal to or greater than the key area presented by the processor (and MSA).

3.2.5.8. SEARCH-KEY-AND-DATA

The three search-key-and-data commands which compare disc unit information to information transferred from the processor are as follows:

- Search-key-and-data-equal
- Search-key-and-data-high
- Search-key-and-data-equal-or-high

The search-key-and-data commands initiate the transfer and comparison of both key and data areas of a record. Bytes may be selectively compared by use of an inhibit byte or bytes (a byte of all binary 1's). If inhibit bytes appear in the information coming from the processor and MSA, then comparison is not made between them and the corresponding disc bytes, because in effect they are masked out by the inhibit bytes. The MSA can search a maximum of 12 bytes. These search commands can be reduced in effect to search-data commands by making the number of inhibit bytes coming from the processor and MSA equal to the number of bytes in the key area on the disc. There are no restrictions on the number or placement of the inhibit bytes. Another disc revolution is needed after a find is made to read the record on which the find was made. These commands may also be used to search (compare) for data areas when there is no key area.

3.2.5.9. CONTINUE-SCAN

Continue-scan commands are used to assist in continuing a search-key-and-data command that is interrupted by the overflow incomplete condition during an overflow record operation. The continue-scan commands are essentially search commands that operate like a read-data command except that data is searched rather than read. The set-file-mask command affects these commands in the same way that the search-key-and-data commands are affected.

SINGLE-TRACK COMMANDS	COMMAND CODE IN HEXADECIMAL	COMMAND CODE IN BINARY 0 1 2 3 4 5 6 7
Continue-scan-equal	25	0 0 1 0 0 1 0 1
Continue-scan-high	45	0 1 0 0 0 1 0 1
Continue-scan-equal-or-high	65	0 1 1 0 0 1 0 1
Continue-scan-no-compare	55	0 1 0 1 0 1 0 1
Continue-scan-set-compare	75	0 1 1 1 0 1 0 1
Continue-scan-set-compare	35	0 0 1 1 0 1 0 1

MULTITRACK COMMANDS	COMMAND CODE IN HEXADECIMAL	COMMAND CODE IN BINARY 0 1 2 3 4 5 6 7
Continue-scan-equal	A5	1 0 1 0 0 1 0 1
Continue-scan-high	C5	1 1 0 0 0 1 0 1
Continue-scan-equal-or-high	E5	1 1 1 0 0 1 0 1
Continue-scan-no-compare	D5	1 1 0 1 0 1 0 1
Continue-scan-set-compare	F5	1 1 1 1 0 1 0 1
Continue-scan-set-compare	B5	1 0 1 1 0 1 0 1

The effects of the continue-scan commands are as follows:

- Continue-scan-equal command performs a search-data-equal operation starting on the positioned record. If the search is successful, the status modifier bit is presented with channel end and device end at the end of the overflow operation.
- Continue-scan-high command performs a search-data-high operation starting on the positioned record. If the search is successful, the status modifier is presented with channel end and device end at the end of the overflow operation.
- Continue-scan-equal-or-high command performs a search-data-equal-or-high operation starting on the positioned record. If the search is successful, the status modifier is presented with channel end and device end at the end of the overflow operation.
- Continue-scan-no-compare command performs a search-data-inhibit-compare operation starting on the positioned record. Channel end and device end are presented at the end of the overflow operation.
- Continue-scan-set-compare command performs a search-data-inhibit-compare operation starting on the positioned record. The status modifier is presented with channel end and device end at the end of the overflow operation.

3.2.6. Sense Commands

When the unit check bit in the status byte is a 1, additional details concerning the nature of any unusual conditions are available in the form of sense bytes (these are explained farther along in this section). The sense-I/O command initiates the presenting of the sense bytes in sequential order to the processor by way of the MSA. Upon completion of an operation caused by a sense command, channel end and device end statuses are presented to the processor by way of the MSA.

3.2.6.1. SENSE-I/O

This command initiates the transfer of applicable information in the form of sense bytes to the MSA and processor. Up to 36 bytes can be transferred under control of the byte count. The first six bytes are used for normal error recovery, and the remainder are used by the diagnostic maintenance routines.

3.2.7. Miscellaneous Commands

These commands cause some action other than a search, seek, sense, read, or write.

3.2.7.1. SET-FILE-MASK

Most commands to the control unit are part of a command chain. The set-file-mask command affects certain subsequent commands in the chain. Although the set-file-mask can be executed anywhere in the chain, issuance of this command as the first command protects the entire chain. If an attempt is made to issue more than one set-file-mask within a chain, a unit check status, associated with command reject, sense information, results.

Execution of this command causes one byte of data, called a mask byte, to be transmitted from the MSA to the control unit. Normal termination of this command is indicated by the presentation of channel end and device end statuses to the processor by way of the MSA.

The effect of the mask byte on write commands is as follows:

Mask Byte	Writes Permitted
00ZYYXXX	All except write-home-address and record R_0
01ZYYXXX	None
10ZYYXXX	Write-data or write-key-and-data
11ZYYXXX	All

When bit Z (bit 2) of the mask byte is set to 1, a data write command (write-data or write-key- and-data) is permitted after a successful short field equal identifier or a key search within the command chain; that is, it permits format writes after a truncated search.

The effect of the mask byte on seek commands is as follows:

Mask Byte	Seek Permitted
YYX00XXX	All
YYX01XXX	Seek-head and seek-cylinder
YYX10XXX	Seek-head
YYX11XXX	None

The file mask is cleared to all 0's under the following conditions:

- End of Command chain
- Control unit is master cleared
- Processor by way of MSA indicates selective reset
- Processor by way of MSA indicates system reset

If any X bit is a 1 in the mask byte or if more than one set-file-mask command is issued within a command chain, unit check status and command reject (bit 0 of sense byte 0) are presented to the processor by way of the MSA.

3.2.7.2. RECALIBRATE

This command, which causes the selected disc unit to seek head 0 and cylinder 0, is used to establish a reference position. It is necessary to have this command because positioning of the accessor is relative to its previous position. If, after long periods of operation, the accessor overshoots or undershoots the intended address, this command reestablishes a reference point. After the new positioning information is transferred from the MSA and processor, channel end is generated. Device end status is presented to the MSA and processor after completion of accessor motion control that has established the cylinder 0 reference point.

3.2.7.3. NO-OPERATION

This command causes channel end and Device end statuses to be presented to the MSA and processor. The addressed disc unit does not respond to this command. If the no-operation command is chained from a format write command, channel end and device end are presented after the control unit has erased the remainder of the track.

3.2.7.4. SPACE-COUNT

This command, which causes an unreadable count area of a record to be spaced (skipped over), is used to gain access to the key and data areas. Three bytes must be transferred from the processor and MSA to the control unit to specify the key and data lengths.

3.2.7.5. ERASE

This command is used to erase to the end of a track after a track overrun has occurred. The command causes the entire track to be erased by writing 1's from the end of the data area of the record on which the preceding search was satisfied; or the record just written by the write-count-key-and-data. After the initiation of this command, channel end and device end statuses are presented to the MSA and processor when the index marker is reached.

3.2.7.6. TEST-I/O

This command initiates interrogation to ascertain the status of an addressed disc unit. If the addressed disc unit is busy, busy status is returned to the MSA and processor. If the disc unit is not operable, unit check is returned. Otherwise, the status byte contains either of the following:

- All 0 bits to indicate the addressed disc is ready for operation.
- The stored status of the addressed disc from an operation prior to the test-I/O.

3.3. STATUS BYTE

Information about the acceptance of a command, the status of the disc subsystem or individual disc unit, and the performance of the command is presented to the MSA in a status byte. The MSA then converts the status byte into a status word. After the status byte is presented to the MSA, the status byte is cleared.

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

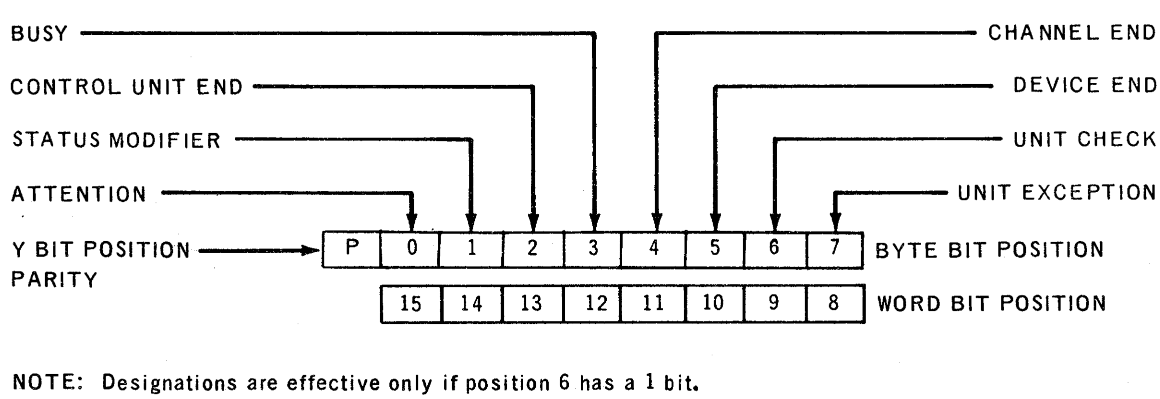

Figure 3–1. Status Byte Format

3.3.1. Attention Bit

This bit is set when the disc subsystem is going into the ready state (power turnon sequence).

3.3.2. Status Modifier Bit

This bit is set as follows:

- when the condition called for in a search command has been met;
- with the busy bit, to indicate that the control unit rather than a disc unit is busy with a previously initiated command;
- parity bit.

3.3.3. Control Unit End Bit

This bit is set when the control unit completes an operation which caused a busy status to be presented to the processor by way of the MSA in response to a request for service.

3.3.4. Busy Bit

The busy bit is set as follows:

- with the status modifier bit to indicate that the control unit is busy;
- without the status modifier bit to indicate that the addressed disc unit is busy;
- with the status modifier bit (in the case of a dual-access configuration) to indicate that the control unit is busy with a previously initiated function on the alternate processor channel.

3.3.5. Channel End Bit

This bit is set when data transfers between the processor and the control unit are completed.

3.3.6. Device End Bit

This bit is set with the channel end bit when any command other than a seek command is completed. When device end is set due to a seek command, it indicates that read/write head motion control has been completed.

3.3.7. Unit Check Bit

This bit is set when a condition results in a 1 appearing in any bits of sense data bytes 0 through 5.

3.3.8. Unit Exception Bit

This bit is set during certain read or write commands where a data length of 0 is detected, and indicates an end-of-file condition. Under this condition, transfers can still take place in the key area but not in the data area. This bit can be set during the following read and write commands:

- Initial-program-load
- Read-key-and-data
- Read-TD-record
- Read-data
- Read-count-key-and-data
- Write-key-and-data
- Write-data

3.4. SENSE DATA BYTES

Sense data bytes are stored data sent to the MSA in response to a sense-I/O command (see 3.2.6). These bytes, which are presented in sequence, give detailed information about unusual conditions detected in the last operation, as well as the current status of the specified disc unit. Sense data bits which are set as a result of an error or fault condition during an operation remain set until cleared by the issuance of a command other than sense or test-I/O.

The bits of the six sense data bytes have the significance shown in Figure 3–2.

SENSE DATA BYTE	BIT 0	BIT 1	BIT 2	BIT 3	BIT 4	BIT 5	BIT 6	BIT 7
0	COMMAND REJECT	INTER-VENTION REQUIRED	BUS OUT CHECK	EQUIP-MENT CHECK	DATA CHECK	OVERRUN	TRACK CONDI-TION CHECK	SEEK CHECK
1	COUNT AREA CHECK	TRACK OVERRUN	CYLIN-DER END	INVALID SEQUENCE	NO RECORD FOUND	FILE PRO-TECTED	MISSING ADDRESS MARKER	OVERFLOW INCOM-PLETE
2	UNSAFE	NOT USED (ALWAYS 0)	NOT USED (ALWAYS 0)	NOT USED (ALWAYS 0)	NOT USED (ALWAYS 0)	UNSE-LECTED STATUS	NOT USED (ALWAYS 0)	NOT USED (ALWAYS 0)
3	READY	ONLINE	UNSAFE	NOT USED (ALWAYS 0)	NOT USED (ALWAYS 0)	END OF CYLIN-DER	NOT USED (ALWAYS 0)	SEEK INCOM-PLETE
4	ALWAYS 0							
5	THIS BYTE IS ALL 0'S EXCEPT WHEN BIT 7 OF BYTE 1 IS SET (OVERFLOW INCOMPLETE). (SEE 3.4.6 FOR DETAILS.)							

Figure 3–2. Sense Data Byte Format

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

  This bit is set as follows:

  – when an invalid command, an invalid sequence of commands, or a command for a feature not installed is received;

  – when the command received is one restricted by set-file-mask;

  – when two set-file-mask commands are sent in the same command chain;

  – when a second disc unit is addressed during a command chain.

- Bit 1 – Intervention Required

  This bit is set when a nonexistent (either physically or electrically) disc unit is addressed.

- Bit 2 – Bus Out Check

  This bit is set when a command or data arrives on the bus out lines with even (incorrect) parity. If the bus out check bit is set, the command reject bit is not set for an invalid command.

- Bit 3 – Equipment Check

  This bit is set to indicate an equipment fault within the subsystem and is set with bit 0 of sense byte 2.

- Bit 4 – Data Check

  This bit is set when an error is detected in the information transferred from a disc unit to the control unit.

- Bit 5 – Overrun

  This bit is set as follows:

  – when the control unit does not receive data bytes within the prescribed time;

  – when data is received too late to be properly written and the remaining record area is filled with 0's;

  – when a subsequent command in a chain is received too late to be properly executed.

- Bit 6 – Track Condition Check

  This bit is set when a read, write, or search command is attempted on a defective track (bit 6 of the flag byte on the record is set if the track is defective). These commands are inhibited for all data records but are permitted for home address and track descriptor records.

- Bit 7 – Seek Check

  This bit is set when an invalid address is sent for a seek command or if less than six address bytes are sent.

3.4.2. Sense Data Byte 1

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

- Bit 0 – Count Area Check

  This bit is set as the result of an error in the count area transferred from a disc unit to the control unit.

- Bit 1 – Track Overrun

  This bit is set when the writing on a track is not completed by the time the index marker is reached.

- Bit 2 – Cylinder End

  This bit is set when a command chain is not completed by the time the end of a cylinder is reached.

- Bit 3 – Invalid Sequence

  This bit is set when two set-file-mask commands are sent in the same command chain. This bit, which is set along with bit 0 (command reject) of sense data byte 0, is also set for an invalid sequence of commands.

- Bit 4 – No Record Found

  This bit can be set only when the M bit of read and search commands is 0 and one of the following conditions exists:

  – Two index markers are detected, and there are no intervening read or write commands during the execution of a chain of search commands.

  – A read or search command has been issued for a blank track.

  – A home address and address marker are missing from a record R_0 track.

- Bit 5 – File Protected

  This bit is set when a seek or write command which has been prohibited by a set-file-mask command is issued.

- Bit 6 – Missing Address Marker

  This bit is set along with bit 4 (Data Check) of Sense Data byte 0 when one of the following conditions exists:

  – Two index markers are passed without detecting any address markers.

  – Two successive records are read in which bit 0 of the flag bytes are equal (indicating that both records are odd or even), and there was no intervening index marker; thus an address marker was missed. A search-ID command inhibits detection of this condition, with the result that is possible to read the remaining data on a track even though an address marker is missing.

- Bit 7 – Overflow Incomplete

  This bit is set when an overflow record is not completed because overflow came either from a defective track or from an alternate track. Bit 6 (defective track check) of sense data byte 0 also is set for these conditions.

3.4.3. Sense Data Byte 2

When the bits of sense data byte 2 are set, they have the following significance:

- Bit 0 – Unsafe

  This bit is set when a disc file malfunction is detected.

- Bit 1

  This bit is not used and is always 0.

- Bit 2 – Missing Clock

  This bit is set when no clocks are received from a device, clocks being the timing (synchronizing) pulses from the device.

- Bit 3

  This bit is not used and is always 0.

- Bit 4 – Priority

  This bit is set when a low-priority command was terminated.

- Bit 5 – Unselected Status

  This bit is set when a file status line is active with no device selected.

- Bit 6 and 7

  These bits are not used and are always 0.

3.4.4. Sense Data Byte 3

When the bits of sense data byte 3 are set, they have the following significance:

- Bit 0 – Ready

  This bit is set when the disc file is ready for operation.

- Bit 1 – Online

  This bit is set when the disc file is online.

- Bit 2 – Unsafe

  This bit is set when a disc file malfunction is detected.

- Bits 3 and 4

  These bits are not used and are always 0.

- Bit 5 – End of Cylinder

  This bit is set when the end of a cylinder is detected.

- Bit 6

  This bit is not used and is always 0.

- Bit 7 – Seek Incomplete

  This bit is set along with bit 7 of sense data byte 1 when a seek command is not successfully completed.

3.4.5. Sense Data Byte 4

The bits of sense data byte 4 are always 0.

3.4.6. Sense Data Byte 5

This byte contains all 0's at all times except when the overflow incomplete bit is set (byte 1, bit 7). The codes in byte 5 indicate the type of command being executed when an overflow incomplete occurs. The codes and their meanings are:

CODE IN HEXADECIMAL	MEANING
06	A read command is in progress.
05	A write command is in progress.
25	A search-key-and-data-equal command is in progress, and the comparison is equal up to this point.
45	A search-key-and-data-high command is in progress, and the comparison is equal up to this point.
65	A search-key-and-data-equal-or-high command is in progress, and the comparison is equal up to this point.
55	Any search-key-and-data operation is in progress, and the comparison is low; or a search-key-and-data-equal is in progress, and the comparison is high.
75	A search-key-and-data-high command or a search-key-and-data-equal-or-high command is in progress, and the comparison is high.

3.5. PROGRAMMING CONSIDERATIONS

This section describes programming and operating conventions pertaining to system reset, device fault indications, and overflow records.

3.5.1. System Reset

The effect produced by the system reset command is the same as that produced by pressing a master reset switch at the processor: the control unit immediately releases control of the interface. No status is presented as a result of a reset.

Any command in progress at the time of the reset is stopped immediately, and any information received as a result of the command should be regarded as invalid.

3.5.2. Device Fault Indicators

The following conditions cause a device check indication on the disc drive control/indicator panel, and results in lockout of the addressed disc drive.

- More than one head selected
- Read and write commands selected simultaneously
- Read and erase commands selected simultaneously
- AC power loss
- DC voltage variation beyond specified limits
- Write current flowing without write gate up
- Erase current flowing without erase gate up
- Write gate up without erase current flowing
- Head selected without seek ready
- Write gate up without AC power
- Write or erase gate up without seek ready

3.5.3. Overflow Records

Considerations for overflow records are as follows:

- Segments other than the first must immediately follow record R_0.
- Segments must not be on a defective track.
- Except for the first and last, all segments must be a complete track in length.

3.5.4. Dual-Channel Access Operation

Dual-channel access is an optional feature. Service to the two channels is first-come, first-served, except as noted in the following.

The channel receiving service first in the case of simultaneously received commands is arbitrarily selected. If the control unit is busy on one channel, a busy bit (3) and status modifier is returned for any commands sent over the other channel. If a unit check status is set at the end of a command, command chaining is inhibited by the channel. For one channel to link a command with its associated sense information, the alternate channel is not allowed access until a sense-I/O command, which is issued by the channel in use, is initiated.

CAUTION

If the last command issued to a device is a seek command which requires motion, or a recalibrate command, the device is reserved for the issuing channel through the same control unit. If the device is addressed through the second channel, a device busy indication will then be returned.

3.6. TIMING

The basic timing components for the disc subsystem are start/stop, accessor movement, and disc rotation time. These times are as follows:

Accessor Positioning Time

one cylinder (minimum) . . . . 20.2 ms

one-third travel (average) . . . . 70.0 ms

full travel (maximum) . . . . 130.0 ms

Disc Rotation

one-half revolution (average). . . . . . 12.5 ms

one revolution (maximum) . . . . 25.0 ms

Head Switching . . . . . 10 µsec

Figure 3–3 shows total access time relative to the number of tracks traveled.

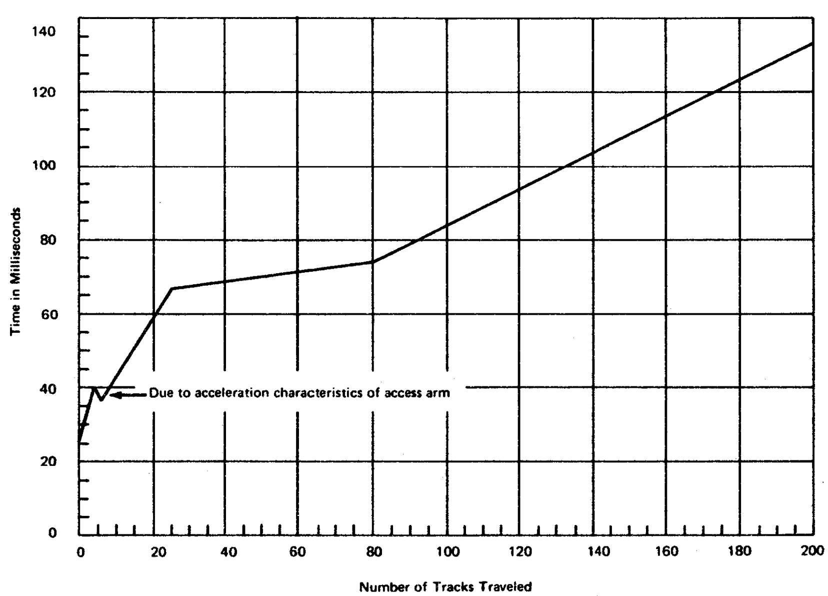

Figure 3–3. Access Times

3.7. MSA WORD FORMATS

The MSA accepts and transfers function words, status words, auxiliary status words and data words in its ooperation. A brief description of these words is provided in the paragraphs which follow. A complete description of these words and the operation of the MSA is found in *UNIVAC 1100 Series Multi-Subsystem Adapter Programmer/Operator Reference, UP-7890* (current version).

3.7.1. Function Word

There are two basic sections in a function word. The least significant 16 bits are the device address and command code for the disc subsystem. The next 12 bits are interpreted by the MSA and specify various formatting and control functions internal to the MSA. Figure 3–4 shows the format of the function word.

B	(shaded)	P	S	R	C	M	X	Q	F	CC	DA	
35	34 28	27 24	23	22	21	20	19	18	17 16	15 8	7 0	BIT POSITION
										0 —— 7	0 —— 7	BYTE POSITION BIT

NOTE: Shaded bit positions are unused and are 0 filled.

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

Byte Information – The notation used on "byte" representation is P01234567 with the least significant bit on the right.

Figure 3–4. Function Word Format

3.7.1.1. B – BOOTSTRAP (BIT 35)

Data format C is automatically selected on bootstrap.

3.7.1.2. P – PARAMETER COUNT (BITS 27 – 24)

This field is used to indicate whether output data is required on a control function when 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. A search truncated is a search-to-equal with the P field equal to 0.

On a jump function, bits 26 through 24 are used to specify the internal address of the next function to be executed in the chain.

3.7.1.3. S – SEARCH FLAG (BIT 23)

This bit is interpreted only on output operations to indicate storage of output (parameter) data in the MSA search identifier register (SIR).

The search flag is ignored unless the SIR is included as an optional feature. The search flag indicates that the specified command is to be reissued to the subsystem, using the SIR as the source of output data. When the first command is executed, the parameter data is simultaneously transmitted to the control unit and loaded in the SIR.

Subsequent commands use the data previously stored in the SIR until the function is completed.

3.7.1.4. R – CHANNEL RESERVE FLAG (BIT 22)

This bit specifies a particular mode of operation when the MSA is accessed by two processor channels. Once a function word or chain of function words have been stored in the MSA, the channel reserve flag controls the channel switching once the 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 the completion of the operation and until a function or chain of functions is successfully completed on that channel in which the channel reserve flag is not specified.

3.7.1.5. C – COMMAND CHAINING FLAG (BIT 21)

This bit specifies that a group of functions are 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.

3.7.1.6. M – MULTIPLE FUNCTION FLAG (BIT 20)

The multiple function flag enables the MSA to execute a string of functions to different devices without storing the 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 the M bit = 0, which defines the last function in a string.

3.7.1.7. X – TRANSLATE (BIT 19)

This bit specifies a translation of output or input data.

Up to two bi-directional translators can be installed in the same MSA module.

Each processor channel is assigned to one MSA translator or the other by way of patch card wiring. Any number of channels can be assigned to a single translator if only one translator is installed. The two available translators provide the following capabilities:

- Fieldata/ASCII Translator

  Processor Fieldata code (64 characters) to Device ASCII Code (64 characters)

  Device ASCII code (96 characters) to Processor Fieldata code (64 characters)

- Fieldata/EBCDIC Translator

  Processor Fieldata code (64 characters) to and from Device EBCDIC Code (64 characters)

3.7.1.8. Q – QUEUING HOLD (BIT 18)

The queuing hold bit 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 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 can also be used on input operations. If the input records do not fill a complete word upon termination of a command, the data is hald 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 on commands requiring output or input data transfer.

3.7.1.9. F – FORMAT SELECT (BITS 17, 16)

These bits specify 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 be in one of the three formats. The following table indicates the format selections:

17	16	Format
0	0	A – Quarter Word
0	1	B – Packed 6-Bit
1	0 or 1	C – Packed 8-Bit

On bootstrap operations, format C is automatically selected. See 3.7.4 for details of the data formats.

3.7.1.10. CC – COMMAND CODE (BITS 15 – 8)

This 8-bit command code is transmitted to the subsystem as the command byte in the initial selection sequence. The general format interpreted by the MSA is as shown in the following:

15 ——— 8	Operation	
XXXX 0000	TEST	
XXXX XX01	WRITE	OUTPUT FORWARD
XXXX XX11	CONTROL	
XXXX 0100	SENSE	INPUT FORWARD
XXXX XX10	READ	
XXXX 1100	INPUT BACKWARD	
XXXX 1000	MSA COMMAND (JUMP)	

3.7.1.11. DA – DEVICE ADDRESS (BITS 7 – 0)

The device address byte specifies the control unit and the disc drive.

O	C	C	C	C	D	D	D	Binary Code
0	1	2	3	4	5	6	7	Device Address Byte
Control Unit					Disc Device			

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

3.7.2. Status Word

The status word is setn to the processor at the completion of each function or chain of functions. The basic word consists of two sections. The least significant 16 bits relate to the disc subsystem status and device address. Bits 16 and 17 refer to error conditions detected in the MSA. The status word format is shown in Figure 3–5.

36–BIT STATUS WORD

(shaded)	ABC	E	DS	DA	
35 – 18	17	16	15 – 8	7 – 0	
			0 – 7	0 – 7	Bytes

ABC – Abnormal Byte Count
E – Error Condition detected by MSA (see Auxiliary Status)
DS – Device Status Byte

ATT	–	Attention	Bit 0
SM	–	Status Modifier	Bit 1
CUE	–	Control Unit End	Bit 2
B	–	Bus	Bit 3
CE	–	Channel End	Bit 4
DE	–	Device End	Bit 5
UC	–	Unit Check	Bit 6
UE	–	Unit Exception	Bit 7

DA – Device Address Byte

NOTE: Shaded Bit positions are unused and should be zero filled.

Figure 3–5. Status Word Format

3.7.3. Auxiliary Status Word

The auxiliary status word contains detailed information relating to error conditions in the MSA. This word is returned in response to a MSA test function as the status word indicating the completion of the operation. The normal status word does not apply in this case. The auxiliary 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 subsystem. The format for the auxiliary status word is shown in Figure 3–6.

(shaded)	NOP	IPE	TCK	CLE	ACP	XCK	LA	SC	WPC	IS	IC	FA	M
35 – 18	17	16	15	14	13	12	11	10	9	8	7	6 – 4	3 – 0

NOTE: Conditions setting bits $2^7 - 2^{17}$ causes bit 2^{16} of the normal Status word to be set. Shaded bit positions are unused and are zerofilled.

NOP – Control Unit Non-Operational
IPE – Input Parity Error
TCK – Time Check
CLE – Control Line Error
ACP – Address Compare Error
} Subsystem Interface Control Checks

XCK – Translate Check
LA – Late Acknowledge
SC – Stall Check
WPC – Word Parity Check

IS – Invalid Sequence
IC – Invalid Command
} Program Checks

FA – Function Address
M – Magnitude of Byte Count

Figure 3–6. Auxiliary Status Word Format

3.7.4. Data Formats

Data formats for word/byte conversion are specified by the function word bits 17 and 16 (see 3.7.1.3). The three formats, A, B, and C, are given in Figure 3–7. The bootstrap function automatically specifies format C.

3.7.4.1. A FORMAT STOP CONTROL

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

3.7.4.2. C FORMAT BIT PADDING AND STRIPPING

When operating in format C, output buffer data may end in the middle of a byte (byte 5 in Figure 3–7). If the control 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 control unit sends the zerofilled byte 5 and presents END status to the MSA. The MSA completes assembly of the word (word 1) and discards the remaining bits (bits 4–7) of byte 5.

FORMAT A (QUARTER WORD)

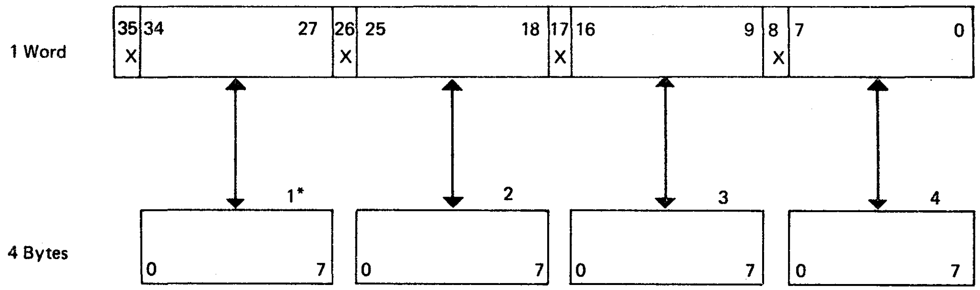

NOTE: Bits 35, 26, 17, and 8 are used for stop control on output operations and forced to binary 0 on input operations.

FORMAT B (6-BIT PACKED)

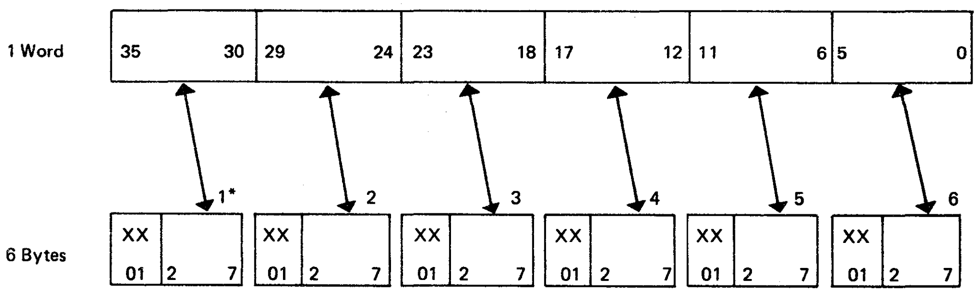

NOTE: Bit 0 and 1 become binary 0 on output and are ignored on input, for each 8-bit byte. When translation is specified, bits 0 and 1 are not forced to binary 0.

FORMAT C (8-BIT PACKED)

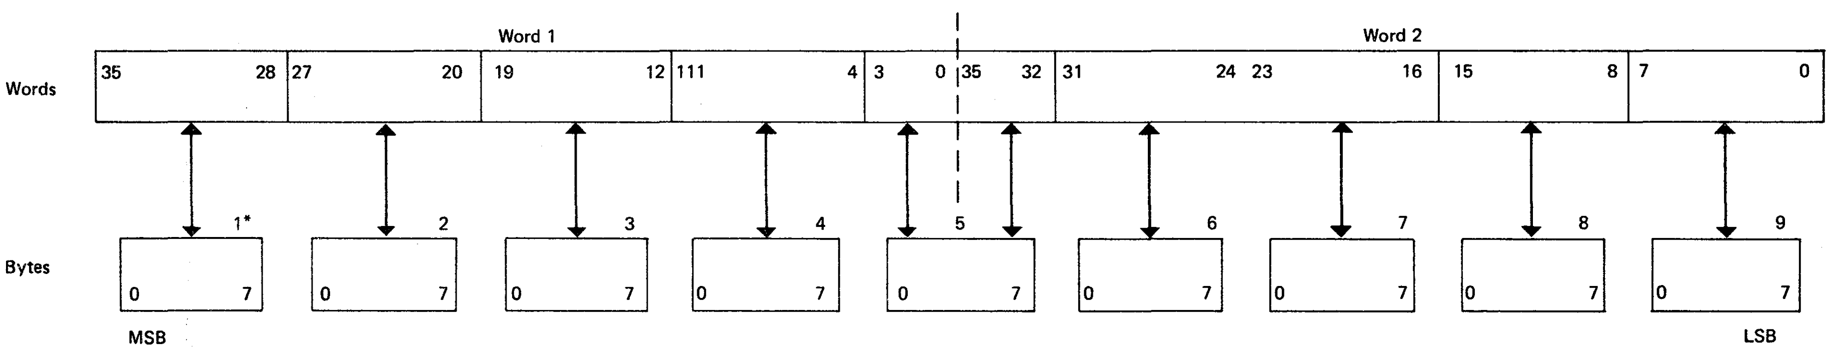

*Numbers on arrows indicate the order of byte transfer.

Figure 3–7. Data Formats

3.7.4.3. OUTPUT DATA TERMINATION

There are three methods of terminating output data transfers to the disc subsystem:

- The stop control bit causes the MSA to terminate data transfers.
- The MSA terminates data transfers because the output buffer and the MSA output queue are both empty and the control unit requests data. The MSA output queue is considered empty if it contains less than one full byte.
- Control unit termination. In this case, the control unit terminates data transfers, by presenting status instead of service in, before or at the same time as the MSA output queue becomes empty.

3.7.4.4. STOP CONTROL

Output operations can be stopped on any full-, half-, or quarter-word boundary when format A is selected. In each data word, certain bit positions which are not used for data bits serve as stop control bits. The first stop control bit terminates data transfers to the control unit. The byte related to the stop control bit (the byte to the right of the stop control bit) is not transmitted to the control unit. If the stop control bit associated with the first byte is 1, no data is transferred to the control unit on the current command.

S C B	BYTE #1	S C B	BYTE #2	S C B	BYTE #3	S C B	BYTE #4
35	34 27	26	25 18	17	16 9	8	7 0

The number of bytes transmitted to the disc subsystem for various settings of the stop control bit is shown as follows:

Word Type	Stop Control Bit Set to 1 in Bit Position	M-Field (Auxiliary Status Word)	ABC Bit (Status Word)	Number of Bytes Transmitted to Disc Subsystem
36 Bit	35 of Word n*	0000	0	4n – 4
	26 of Word n	0001	1	4n – 3
	17 of Word n	0010	1	4n – 2
	8 of Word n	0011	1	4n – 1

n = 0, 1, 2, 3, ...

*n = 0. If n = 0, no data is transferred to the control unit. The ABC bit will then be a 0 but the M-field remains unchanged from the last instruction. In this case, the M-field is not pertinent and is ignored because the ABC bit = 0.

3.7.4.5. LIMITED BUFFER (EMPTY OUTPUT QUEUE)

The second method of terminating data transfers is by specifying a buffer length such that the control unit requests data after the MSA output queue and the processor buffer are both empty. Indications caused by this method of termination are shown as follows:

OUTPUT					
Format	Number of Words Transmitted by Processor	Number of Bytes Transmitted to Disc Subsystem	Number of 0's Padded in Last Byte Transmitted to Disc Subsystem	ABC Bit (Status Word)	M-Field (Auxiliary Status Word)
A	n*	4n	0	0	0000
B	n*	6n	0	0	0000
C	2n*	9n	0	0	0000
	2n+1	9n+5	4	1	0101

n = 0, 1, 2, 3, ...

*n = 0. If n = 0 the command is not executed. If it is a single output function or an output function within a chain, the MSA does not initiate the function or chain of functions and stall check is indicated. If the output function is detected during the execution of a chain, the function is not executed and the function address (FA) field indicates the address of the last successful command.

3.7.4.6. CONTROL UNIT TERMINATION (OUTPUT)

The final method of termination is caused by the control unit terminating data transfers (by presenting status in to the MSA) before the empty output queue condition occurs. This termination is described in chart form as follows:

OUTPUT — 36-Bit Word					
Format	Number of Bytes Transmitted to Disc Subsystem	ABC Bit (Status Word)	M-Field (Auxiliary Status Word)	Number of Words† Transmitted by Processor	Number of 0's† Padded in Last Byte Transmitted to Disc Subsystem
A	4n*	0	0000	n	0
	4n+1	1	0001	n+1	0
	4n+2	1	0010	n+1	0
	4n+3	1	0011	n+1	0
B	6n*	0	0000	n	0
	6n+1	1	0001	n+1	0
	6n+2	1	0010	n+1	0
	6n+3	1	0011	n+1	0
	6n+4	1	0100	n+1	0
	6n+5	1	0101	n+1	0
C	9n*	0	0000	2n	0
	9n+1	1	0001	2n+1	0
	9n+2	1	0010	2n+1	0
	9n+3	1	0011	2n+1	0
	9n+4	1	0100	2n+1	0
	9n+5	1	0110	2n+1	4
	9n+6	1	0101	2n+2	0
	9n+7	1	0111	2n+2	0
	9n+8	1	1000	2n+2	0

n = 0, 1, 2, 3, ...

*n ≠ 0. If n = 0, the contents of the M-field remains unchanged from the previous command. In this case, the M-field is not pertinent and should be ignored because the ABC Bit = 0.

† = not a consideration if transfer is halted by the control unit.

3.7.4.7. INPUT TERMINATION

There are two methods of normal termination for input operations. The first method is by disc control unit termination. This occurs when the disc control unit sends all the data it has available and the processor accepts all this data. The terminating indications for this case are shown in the chart that follows.

The second method is the result of the processor buffer being limited such that not all the data transmitted by the control unit are accepted by the processor by way of the MSA. In this case, the ABC bit is set to one and the M-field in the auxiliary status word becomes 0000 regardless of any coincidence of word and byte boundaries. The data remaining in the MSA input queue is discarded.

INPUT — 36-Bit Word					
Format	Number of Bytes Received by MSA from Control Unit	M-Field (Auxiliary Status Word)	ABC Bit (Status Word)	Valid Data Bits in Last Word Sent to Processor	
				Input Forward	Input Backward
A	4n	0000	0	34–27; 25–18; 16–09; 07–00	34–27; 25–18; 16–09; 07–00
	4n+1	0001	1	34–27	07–00
	4n+2	0010	1	34–27; 25–18	16–09; 07–00
	4n+3	0011	11	34–27; 25–18; 16–09	25–18; 16–09; 07–00
B	6n	0000	0	ALL	ALL
	6n+1	0001	1	35–30	05–00
	6n+2	0010	1	35–24	11–00
	6n+3	0011	1	35–18	17–00
	6n+4	0100	1	35–12	23–00
	6n+5	0101	1	35–06	29–00
C	9n	0000	0	ALL	ALL
	9n+1	0001	1	35–28	07–00
	9n+2	0010	1	35–20	15–00
	9n+3	0011	1	35–12	23–00
	9n+4	0100	1	35–04	31–00
	9n+5	0101	1	ALL*	ALL*
	9n+6	0110	1	35–24	11–00
	9n+7	0111	1	35–16	19–00
	9n+8	1000	1	35–08	27–00

n = 0, 1, 2, 3, ...

*Bit Stripping Occurs

3.8. MSA FUNCTIONS

Two functions can be executed by the MSA without issuing commands to the disc subsystem. MSA functions are specified by the MSA device address $F1_{16}$.

3.8.1. Test

This function is used to return detailed status information to the processor. Once issued to the MSA, the test function causes the auxiliary status word to be returned by way of an external interrupt. To select the MSA, the DA field of the function word should be $F1_{16}$. The test function should be used when a status word is returned indicating a MSA error. 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 complete. Status words are generated in response to this MSA test and nontest functions.

3.8.2. Jump

This function is used to repeat a chain of functions or a section of the chain without interrupting the processor. When utilized, it must always be the last function in the chain.

3.8.3. Nontest Functions

Functions other than those specifying test commands cause an external interrupt at the disconnect point; that is, when the MSA disconnects from the disc subsystem. This point is defined as the acceptance of the first status byte after initial selection sequence or the status byte containing the channel end bit, whichever occurs first. The status byte accepted at the disconnect point is returned in the status word. Status is not accumulated in the MSA.

3.8.3.1. CONTROL UNIT BUSY

If a control unit is busy or has status pending for a device other than that specified in the DA field of the function word, the MSA responds with the status modifier and busy bits indicated in the status word.

3.8.3.2. PENDING STATUS

When the control unit has a status byte stored for the addressed device, a nontest function is rejected. The MSA generates an external interrupt at the completion of initial selection sequence. The status word indicates the status stored in the control unit along with the busy bit.

3.9. DELAYED INTERRUPTS

Seek commands which result in physical motion of the head boom generate two interrupts. The first is immediate unless eliminated by use of the M (multiple function) flag. The second is generated when the head boom reaches its destination. This interrupt is submitted to the processor when the command or chain of commands in progress have been completed and their interrupt acknowledged. The processor, therefore, must be capable of accepting a data interrupt plus seven consecutive seek complete interrupts.

3.10. DETAILED DEVICE STATUS (SENSE BYTES)

When the regular status word indicates unit check (bit 9), additional details of the device status may be solicited by issuing a sense command for that device. The sense command stores the current status of the selected disc drive, and then transfers one to six sense bytes to the processor channel by way of the MSA. The information in the sense bytes contains the current status of the disc drive and any error or fault conditions that occurred during the last operation. The regular status byte presented at the completion of the sense command contains channel end (bit 4) and device end (bit 5). Any nonready condition in the selected disc drive is not indicated in the status byte. See 3.4.

In the MSA, the sense bytes are treated as input data and go through the data conversion and formatting. After conversion these sense bytes are sent to the processor as input data words.

4. PROGRAMMING EXAMPLE

The following listing is an example of programming for a disc subsystem to position two disc files (disc 0 and disc 1) and a continuous write on disc 4 using A format. Refer to 3.7.1 – 3.7.1.10. Each line represents bits 27 through 0 of a UNIVAC 1108 processor word.

	P	S	R	C	M	X	Q	F	Command	Device Address
	27——24	23	22	21	20	19	18	17 16	15——8	7 ——0
External Function	1000	0		0	1	0			Seek	0 0
Parameter 1										
Parameter 2										
External Function	1000	0		0	1	0			Seek	01
Parameter 1										
Parameter 2										
External Function	1000	0		1	0	0			Set-file-mask	0100
External Function	1000	0		0	1	0			Seek*	01
External Function	0101	1		1	0	0			Search-ID-equal	**0000
External Function	1000	1		1	0				Write	**0000
External Function	0000	1		1	0	0			Search-ID-truncated	**0000
External Function	0010	0		0	0				Jump	1 —— 1
Mask Byte Parameter										
Search Parameter 1										
Search Parameter 2										
Search Parameter 3										
Search Parameter 4										
Search Parameter 5										

* It is advisable to place a seek between the set-file-mask and search to ensure the access arem is on position. Also, without the seek, it is not possible to have the M-bit in the command. Without the seek in this case, the control unit will give a unit check and disconnect.

** No address is necessary in these commands. Prior issuance of the set-file-mask protects these commands.

The first function of a group sent by the processor is a forced function. The remainder of the group can be sent as buffered functions or data. In this example, the M bit is set (bit 20). The M bit enables the MSA to execute a seek function to different devices without interrupting the processor twice on each function. If the head boom requires movement, channel end (seek initiated) is discarded by the MSA and the next function in the string performed. However, if the head boom is in place or an error is detected, the string is broken and an interrupt presented. The MSA keeps calling out functions until the M bit=0, which identifies the last function in a string. After initiation of the "seek device 0", channel end status is discarded by the MSA. The same status is then discarded when a seek is initiated at address 1.

In the next function received from the processor (set-file-mask), there is a 1 in the C field, bit 21; and in bit position 27 of P. The C bit indicates that once the initial function word is stored in the MSA, further function words are to be requested and stored in the MSA function buffer register. Bit 27 in P indicates that output data is required by this control command after it is transferred to the disc control.

The next function word following a set-file-mask and seek is a search-ID-equal, which has the search flag bit set. This indicates that the specified command is to be reissued to the disc control using the data stored in the PSR as the source of the output data.

The 5_{16} indicated in the P position indicates that five words of parameters are to be stored in the SIR (search identifier) register. This becomes the output data issued to the disc subsystem with the information needed for search-ID.

The write instruction is automatically loaded into the function register. This function was requested because the C (chain) bit was set in the search-ID-equal function.

The search-ID-truncated (same code as a search) is stored in the function register. This function is requested because the C bit is set into the previous function word.

The jump instruction is stored in a function register also. In this particular function, the C bit is inactive indicating no request for more information. The jump address is indicated by the lower three bits in P and points in the example to the third function register (the write function).

Once the chain is stored and it is determined by the decoding of the function word that output data is required, ODR (output data request) is activated. When the ODR is acknowledged (data available to be written), the MSA releases the set-file-mask function to the control unit. The parameter word for this function contains one byte of information. The third most significant bit of this byte must be set to 1 to allow a write after truncated search.

Upon completion of the set-file-mask function the search-ID-equal function is released to the control unit. After completion of the search-ID-equal function by way of the status between the control unit and MSA, the write function is then released to the control unit by the MSA. The MSA attempts to keep its output Q registers full by presenting data requests to the processor. After writing a sector status, the control unit instructs the MSA to initiate the truncated search, which is a means of forcing multiple sector operation on the control unit without breaking up the output data stream. The truncated search makes a match and the jump function is activated from the function register. On a jump function, the last three significant bits of the P field are used to specify the internal address of the next function to be executed. This will then be a write function. These functions are repeated until the processor detects an end of buffer and stops sending data. The status received from the control unit to the MSA indicates a successful completion. The status the processor receives is bits 10 and 11, channel end and device end, bits 8, 9, 12 – 16 should be 0, bit 17 may be 0 or 1 depending on the output data length and format.

APPENDIX A. UNIVAC 8411 DISC SUBSYSTEM

CHARACTERISTICS	
Disc Drum Per Subsystem	1 to 8
Disc Drives Per Cabinet	1
Read/Write Head	1
Read/Write Heads per Disc Drive	10
Discs Per Pack	6
Recording Surfaces Per Disc Drive	10
Tracks Per Disc Surface	203
Maximum Eight-Bit Bytes Per Track	3,025
Data Capacity (Eight-Bit Bytes) Per Disc Pack	7,250,000
Track Density	100 per inch
Bit Density	1100 bits per inch
Data Transfer Rate	156,000 bytes per second
Disc Speed	2400 rpm
Disc Diameter	14 inches
Vertical Spacing Between Discs	0.350 inch
Parity (Over Interface Lines)	Add byte parity on command, address, status, sense, and data bytes; longitudinal parity on all record fields.

Table A–1. Characteristics of UNIVAC 8411 Disc Subsystem

Component/Optional Feature	Type Number		Requirements for Subsystem
	60 Hz	50 Hz	
Control Unit	5024-00	5024-01	1
Dual Channel	F1043-00	F1043-00	1 per control unit, optional
Record Overflow	F1098-00	F1098-00	1 per control unit, optional
File Scan	F1099-00	F1099-00	1 per control unit, optional
Disc Drive	8411-00	8411-01	1 minimum 8 maximum
Disc Pack	F1211-00	F1211-00	1 per disc drive minimum

Table A–2. Basic Components and Optional Features for 8411 Disc Subsystem

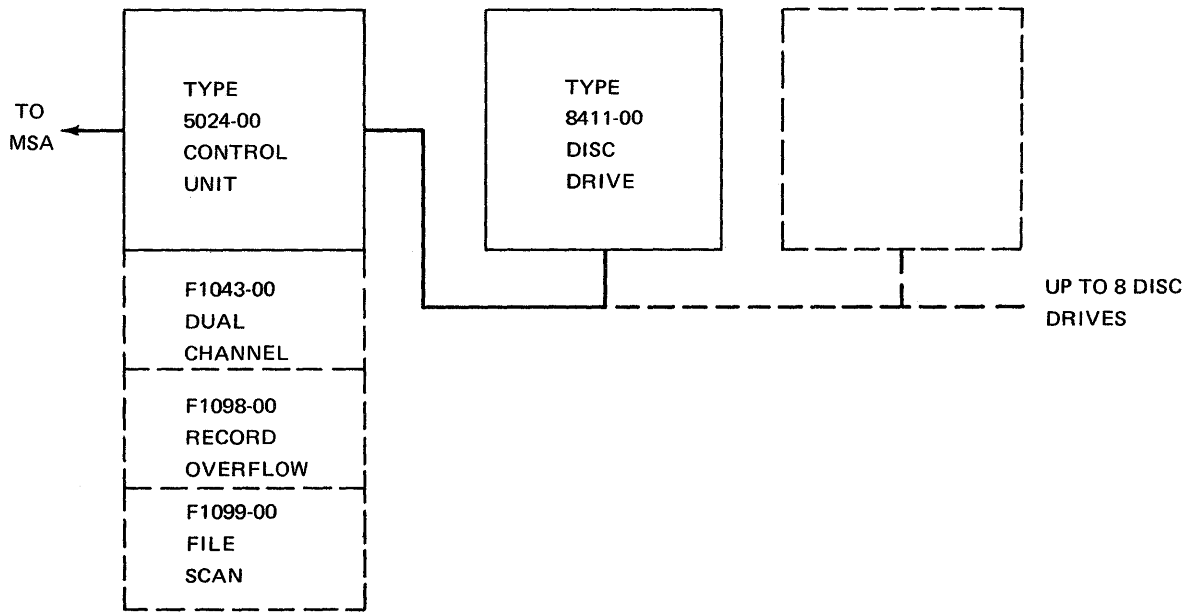

Figure A–1. Operating Configuration for the UNIVAC 8411 Disc Subsystem

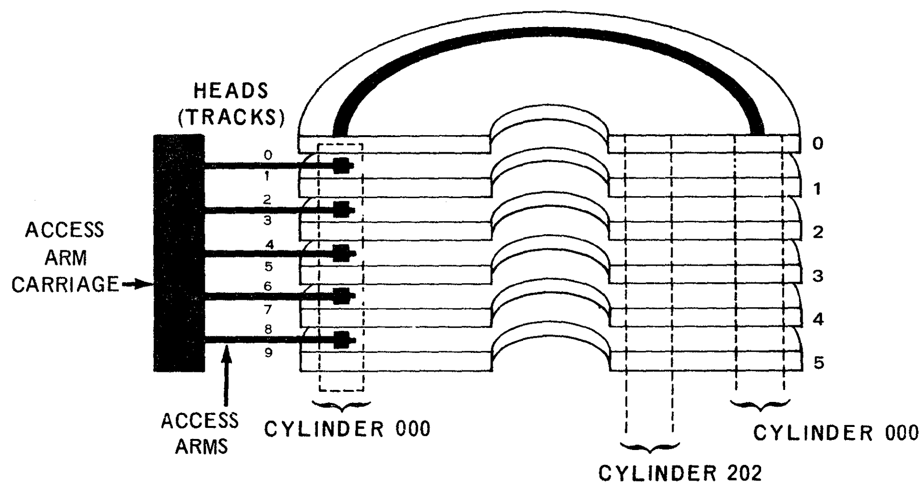

	U	C1	C2	H1	H2
BIT	01234567	01234567	01234567	01234567	01234567
	BYTE 1	BYTE 2	BYTE 3	BYTE 4	BYTE 5

Legend:

U = disc drive unit address (unit is identified by bits 5, 6, and 7, which are used to represent 0 through 7 in binary notation); the control unit is identified by bits 0 through 4.

C1 and C2 = cylinder address; C1 must always be 0; C2 is the cylinder identification (from 000 through 202 in binary notation).

H1 and H2 = head address; H1 must always be 0; H2 is the track digit (from 0 through 9 in binary notation); bit position 7 of byte 5 is the least significant bit. It should be noted that the words head and track are synonomous.

Figure A–2. UNIVAC 8411 Disc Subsystem File Organization and Addressing Format

MAXIMUM NUMBER OF DATA BYTES PER RECORD	NUMBER OF EQUAL-LENGTH RECORDS PER TRACK																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
WITHOUT KEY FIELD	3625	1740	1131	830	651	532	447	384	334	295	263	236	213	193	177	162	149	138	127	118
WITH KEY FIELD	3605	1720	1111	811	632	512	428	364	315	275	244	217	194	174	158	143	130	119	108	99

Table A–3. UNIVAC 8411 Direct Access Subsystem Bytes Per Record

APPENDIX B. FORMAT GAP BIT CONFIGURATION

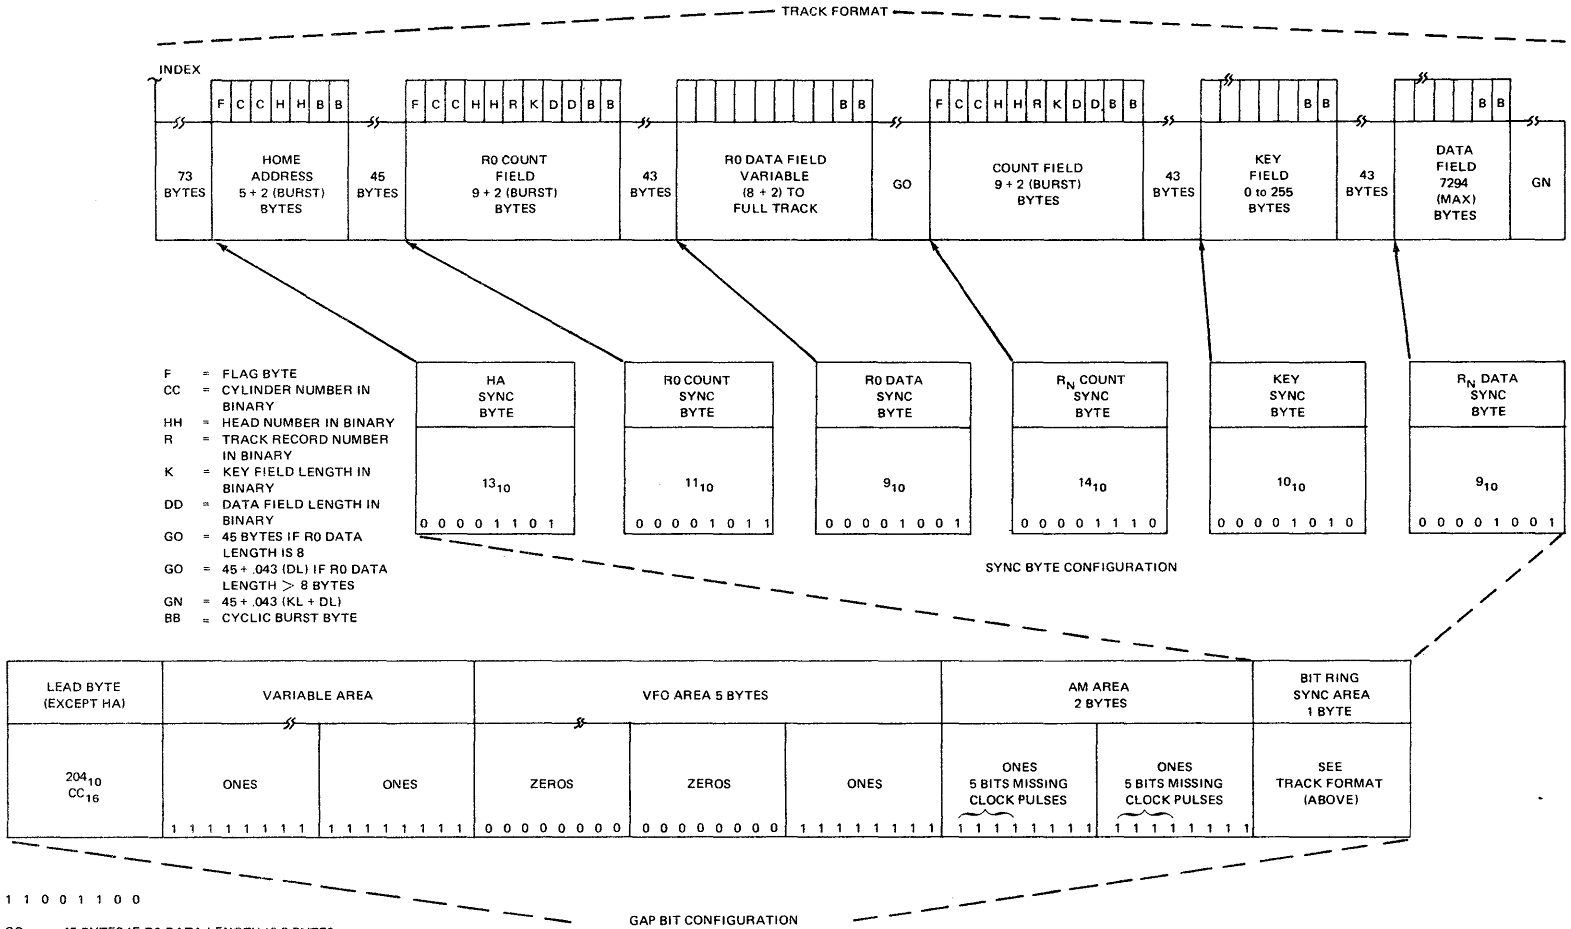

1 1 0 0 1 1 0 0

GO = 45 BYTES IF R0 DATA LENGTH IS 8 BYTES
GO = 45 + .043 (DL) IF R0 DATA LENGTH > 8 BYTES
GN = 45 + .043 (KL + DL)

APPENDIX C. TRACK ORGANIZATION