

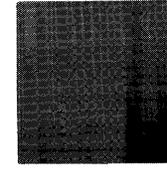
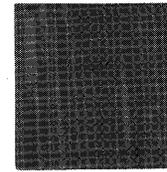
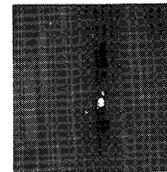
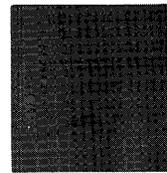
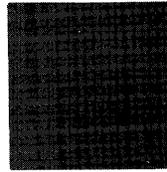
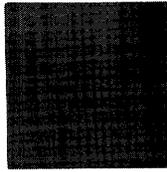
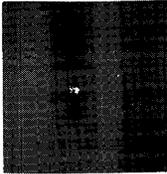
## Systems Reference Library

### IBM 1311 Disk Storage Drive

This publication describes the operation of the IBM 1311 Disk Storage Drive.

The functional and operating characteristics of the drive are presented. The timing considerations related to the drives are shown.

The special features that are available on the 1311 Disk Storage Drive are described, along with the instructions for each feature when applicable.



This publication, A24-3086-3, is a reprint of A24-3086-2 except for the removal of the Pack-On Light description and addition of an Enable-Disable switch description.

Copies of this and other IBM publications can be obtained through IBM Branch Offices. Comments concerning the contents of this publication may be addressed to:  
IBM, Product Publications Department, San Jose, Calif. 95114

## Contents

<b>IBM 1311 Disk Storage Drive</b> .....	5
Model 1 .....	5
Model 2 .....	5
Model 3 .....	5
Model 4 .....	5
Model 5 .....	5
Storage Capacity .....	6
Access Assembly .....	6
Access Speed .....	6
Disk-Storage Organization .....	7
Cylinder Concept .....	7
Disk-Storage Addressing .....	7
<b>Disk-Storage Drive Keys and Lights</b> .....	9
Start-Stop Key-Light .....	9
Write Address Key-Light (Drive 0 only) .....	9
Module Ready Light .....	9
Select Lock Light (Disk-Storage Drive 0 only) .....	9
Enable-Disable Switch (Meter) .....	9
Compare Disable Switch (C.E.) .....	9
<b>Special Features</b> .....	10
Track Record Feature (Models 4 and 5 only) .....	10
Read Disk-Track Record .....	10
Write Disk-Track Record .....	11
Read Disk-Track Record with Address .....	11
Write Disk-Track Record with Address .....	12
Seek Overlap (Models 4 and 5 only) .....	13
Scan Disk Feature .....	13
Direct Seek Feature (Model 4 only, Standard on Model 5) .....	14
Disk-Control Field .....	14
Direct Seek (8 2 1) # .....	15
Direct Seek Timing .....	15
Disk-Storage Timing .....	15
Seeking Disk-Storage Records .....	15
Dummy Seek to Cylinder 00 .....	16
Timing Considerations for Reading and Writing .....	16



Figure 1. IBM 1311 Disk Storage Drive

## IBM 1311 Disk Storage Drive

The IBM 1311 Disk Storage Drive (Figure 1) provides a fast, efficient means of disk storage for any system to which it is attached. Each drive is capable of storing 2 to 2.98 million alphanumerical characters.

There are five models of the IBM 1311 Disk Storage Drive. Various combinations of these models are used on the systems that can attach the 1311.

### Model 1

The 1311, Model 1, is the required first drive when 1311's are attached to either an IBM 1440 Data Processing System or an IBM 1460 Data Processing System. This model is the master drive on the system and contains controlling circuitry, lights and switches that control it, and other 1311's attached to the system. Only one Model 1 can be attached to a system and four additional Model 2 drives.

### Model 2

The 1311, Model 2, is the disk storage drive used when more drives than the master are required on the system. The Model 2 drives are attached to a system through the master drive, and a maximum of four Model 2 drives can be attached. The Model 2 drives can attach to these systems through the appropriate master drive:

- IBM 1620 Data Processing System
- IBM 1710 Control System
- IBM 1401 Data Processing System
- IBM 1410 Data Processing System
- IBM 1440 Data Processing System
- IBM 1460 Data Processing System

### Model 3

The 1311, Model 3, is the required first drive when 1311's are attached to either an IBM 1620 Data Processing System or an IBM 1710 Control System. This model is the master drive on the system and contains power and controlling circuitry, lights and switches that control it, and other 1311's attached to that system. Only one Model 3 can be attached to the system and three additional units, Model 2, for a maximum of four disk-storage drives on a 1620/1710 system.

### Model 4

The 1311, Model 4, is the required first drive when 1311's are attached to the IBM 1401 Data Processing System. This model is the master drive on the system

and contains the power supply and the controlling circuitry, lights and switches that control it, and the other 1311's attached to that system. Only one Model 4 can be attached to the system and four additional units, Model 2, for a maximum of five disk-storage drives on a 1401 system.

### Model 5

The 1311, Model 5, is the required first drive when 1311's are attached to the IBM 1410 Data Processing System. This model is the master drive on the system and contains the power supply and the controlling circuitry, lights and switches that control it. Only one Model 5 can be attached to a channel and four additional units, Model 2, for a maximum of five disk-storage drives on a channel.

Information is stored on small removable disk packs. This concept of removable disk packs means that only those records needed for a particular application need be used. Data records for other applications can be removed from the system and stored.

The IBM 1311 Disk Storage Drive and the portable disk pack allow the user to select the number of disk drives (maximum of 5) needed on the system at any one time and to place only the required disk records on the drives. Any number of disk packs are available for data storage. The number of disk packs needed by the user depends on his disk-storage requirements.

The disk pack is a precision device, designed to meet the extremely close tolerances required for high-speed reading and writing on the 1311 Disk Storage Drive. When the disk pack is not on a disk-storage drive, it is inside its protective cover (Figure 2). The disk pack and its cover combine to make a rugged, sealed container. Because the disk pack is a precision instrument, it should be handled with care. Every effort should be made to keep dust and foreign matter away from the disk surfaces. (The dust could become wedged between the disk surface and the read-write heads, causing permanent damage to the head and the disk.)

The disk drives used in combination with any one of the IBM Data Processing Systems (1401, 1410, 1440, or 1620/1710) expands the input/output abilities of that system and provides the user with a flexible, low-cost in-line data processing system.

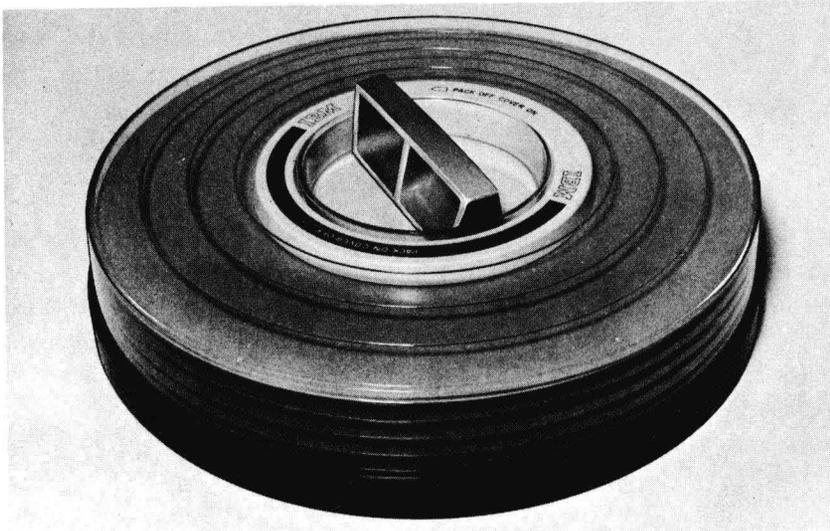


Figure 2. IBM 1311 Disk Storage Drive Disk Pack

### Storage Capacity

The IBM Data Processing System (1401, 1410, 1440) can have as many as five IBM 1311 Disk Storage Drives attached to it. The 1620/1710 system can have as many as four drives attached to it. Each drive has its own independent access mechanism consisting of ten read-write heads. All drives can be used by a single program in much the same way a tape system uses tape units for storing master and active input and output files. The 1311 supplies only the access and drive mechanisms; the disk records are separate.

Independent portable disk packs are used with the 1311 drive. Each disk pack consists of six magnetic disks, 14 inches in diameter, providing ten surfaces for recording data. The pack is easily removed (nominal change time of two minutes).

The packs are interchangeable with other IBM 1311 drives. This means that the philosophy of tape libraries can now be applied to magnetic-disk storage. A library of disk packs can be set up, and when needed, the packs necessary for an operation can be placed on the disk-storage drives for system use.

Each disk pack has a storage capacity of 2 million alphameric characters when records are written in the sector mode, or 2.98 million when the data is written in the track-record mode. With five disk drives attached to a system, from 10 to approximately 15 million alphameric characters of data are available for system use.

### Access Assembly

The access assembly on each IBM 1311 disk storage drive has 10 read-write heads attached to it (Figure 3).

The read-write heads move horizontally across the tracks on the disk during a seek operation. Vertical movement of the access mechanism is not necessary,

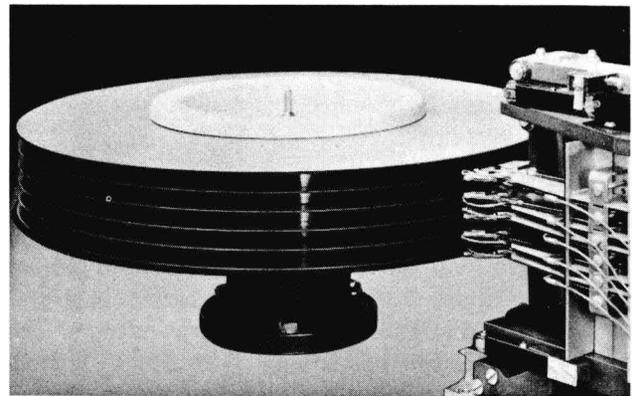


Figure 3. Access Assembly

because there is a head for each disk surface. Because each disk surface has its own read-write head, it is not necessary to perform a seek operation if the sought record is in the same vertical plane (cylinder) as the record previously operated upon.

### Access Speed

The use of a comb-like access assembly greatly reduces the access time of a record in disk storage. The maximum access time is 400 ms; the average access time is 250 ms. Access time can be further reduced by using Direct Seek (special feature), which does not require the access assembly to return to the home position for each seek operation. The Direct Seek feature reduces the maximum access time to 250 ms and reduces the average access time to 150 ms.

Another timing factor that must be considered is rotational delay time. Rotational delay is the time required for the starting point of the disk record sought to come under the read-write head. The maximum time

for rotational delay on an IBM 1311 disk drive is 42 ms; the minimum time is 2 ms (time required to turn on the read-write head); and the average time is 22 ms. Therefore, when considering the time required to make a disk record available to the system, consider both access-motion time and rotational-delay time.

### Disk-Storage Organization

The ten disk surfaces are divided into tracks and sectors similar to the previous disk-storage devices.

Each disk surface contains one hundred tracks, which are subdivided into twenty sectors per track (Figure 4). Thus there are twenty thousand addressable sectors per disk pack. The capacity of a sector depends upon the mode in which the writing is executed. Sectors written in the load mode contain 90 data characters. Sectors written in the move mode contain 100 data characters.

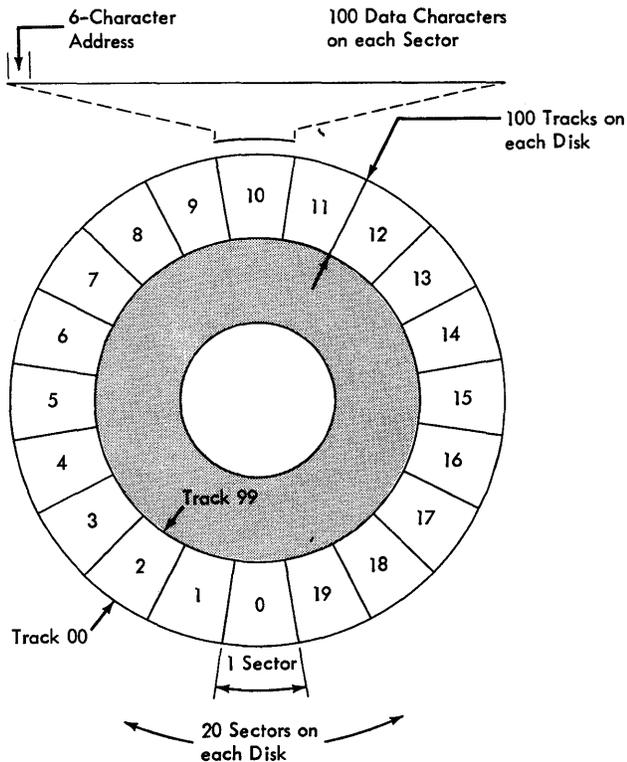


Figure 4. Disk Sector Organization

### Cylinder Concept

The combination of the ten disk surfaces in a parallel plane and the multiple-head access assembly constitute a powerful disk organization. The ten read-write heads of the access assembly are at any one time located in the same relative position on each of the parallel records. Thus, ten tracks in the same vertical plane

are available with one seek operation (Figure 5). Therefore, this vertical plane can be considered a cylinder with ten tracks (200 sectors) of data available for system use with a minimum of time required for access to them.

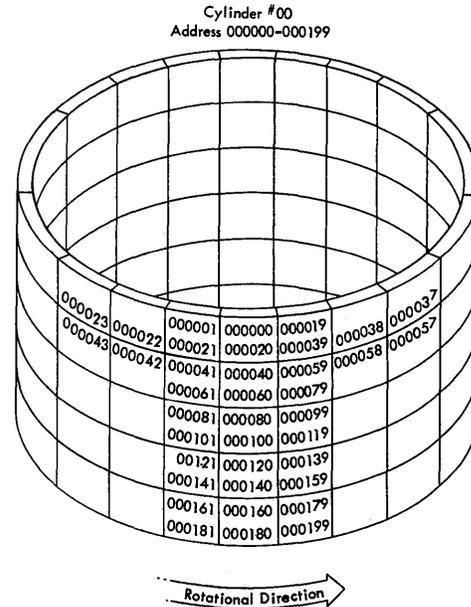


Figure 5. Cylinder Concept

Because each disk has 100 tracks per surface, consider that each 1311 disk drive contains 100 cylinders (Figure 6). Data within a cylinder can be reached in a minimum amount of time, that is, rotational-delay time only. There is no need to move the access assembly, because all the read-write heads are at the tracks that comprise the cylinder.

The number of characters in a cylinder can vary. If the records are written in the sector mode, each cylinder contains 20,000 characters. If the records are written in the track-record mode, each cylinder contains 29,800 characters.

Sector addresses can be repeated on a disk track. This makes possible another technique for data storage, called the *revolver method*. The same data is stored in several records on the same track, thus reducing the rotational time required for access to the data. This method is used for the storage of tables, constants, and other reference data that is frequently required for system use.

### Disk-Storage Addressing

A six-digit sector address precedes each addressable location in a disk pack. These numeric addresses are normally sequential within track, cylinder, and disk pack. If a system uses five drives, the sector addresses continue in this sequence from disk drive to disk drive.

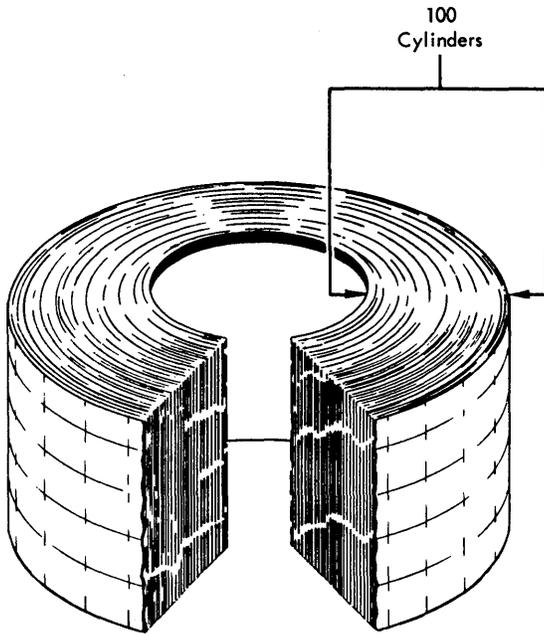


Figure 6. Number of Cylinders

Figure 7 is an illustration of the sequence of addresses in a disk pack.

The address ranges for additional disk drives are shown in Figure 8.

To find the physical location of a disk record, disregard the first digit of the six-digit sector address and divide the address remainder by twenty. The quotient is made up of the drive, cylinder, and track numbers; the remainder is the sector number (Figure 9).

The sector addresses recorded on the disk pack are compared to the core-sector-address portion of the 10-digit disk-control field (digits 2-7). (Digit one of the disk-control field contains the alternate control digit or the disk number; digits 8-9 and 10 of the disk-control field contain the number of sectors to be operated on.) The disk-storage instruction B-address is the core-storage address that contains the disk-drive number or the alternate control-digit portion of the disk-control field.

Sector Address	LOCATION		
	Cylinder	Track	Sector
000000	1st	1st	1st
000019	1st	1st	last
000199	1st	last	last
019999	last	last	last

Figure 7. Sequence Addresses

Disk Drive Number	Address Range
0	000000-019999
2	020000-039999
4	040000-059999
6	060000-079999
8	080000-099999

Figure 8. Disk-Drive Address Ranges

Sector Address 015734

	cylinder	
drive		track
	┌	└
	0 7 8 6	
20	15734	
	┌	
	140	
	└	
	173	
	┌	
	160	
	└	
	134	
	┌	
	120	
	└	
	14	
	└	
	sector	

Therefore, sector address 015734 is in the first drive cylinder number 78 on sector number 14 of track number 6.

Figure 9. Physical Location of a Sector Address

## Disk Storage Drive Keys and Lights

### Start-Stop Key-Light

Pressing this key supplies the power to the disk drive and establishes the unit in ready condition (Figure 10). The key light is turned on.

Pressing this key again retracts the access mechanism to the home position and removes the power from the drive. With the unit in this state, disk-pack replacement can be accomplished without affecting the operation of the system or other drives under control of the attached system program.

### Write Address Key-Light (Disk-Storage Drive 0 only)

This key is located on the master disk drive (number 0) and controls the write disk address mode for all disk drives. Write operations in the address mode can be performed only when this key-light is pressed and the light is on. The WRITE DISK SECTOR and the WRITE DISK TRACK RECORD instructions (without addresses) cannot be performed if the light is ON.

When this key-light is OFF, *write with address* operations cannot be performed. Read operations can be performed with or without addresses, regardless of the setting of this key-light.

If the write operation called for by the stored program does not agree with the key-light setting, the instruction is terminated and a program testable indicator is set.

### Module Ready Light

This light is ON when the IBM 1311 is ready for use by the attached system.

The IBM 1311 is in ready state when the following conditions have been fulfilled.

- The disk pack is rotating up to speed
- The access mechanism is extended in to the disk pack
- The read-write heads are loaded.

This light is OFF when the power is removed from the drive. If it is OFF, processing continues until the

stored program addresses this particular drive; then the instruction terminates and a program testable indicator is set.

### Select Lock Light (Disk-Storage Drive 0 only)

This light comes ON to signal that the disk drives on the system are inoperable and customer engineering assistance is needed.

### Enable-Disable Switch (Meter)

The Enable-Disable switch must be in the Enable position in order to operate the 1311 master drive. When the switch is in the Enable position, the 1311 meter will run provided the CPU meter is also running. The position of the switch may be changed while the CPU is running but the status of the 1311 will not change until the CPU has stopped.

If satellite files are installed, the satellites are interlocked with the master drive so that the satellites cannot be used unless the Enable-Disable switches on both master drive and satellites are in the Enable position.

The satellite meter will stop while a disk pack is being changed. The master drive meter will stop during disk pack change only if no satellite drives are selected for meter operation.

### Compare Disable Switch (C.E.)

This switch is located inside the master drive (number 0). It is to be used only by the customer engineer. It is a pull-out-type switch and is restored to normal position when the rear gate is closed after the switch has been used.

When the switch is pulled out, the sector address compare during the address mode of track read operation is disabled. This permits the reading of track information when an unequal address compare error prevents reading from the record.

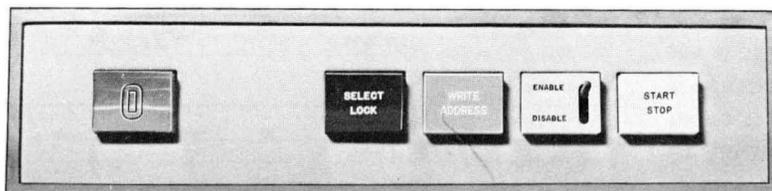


Figure 10. IBM 1311 Control Keys and Lights (Master Unit)

## Special Features

The special features provided for the IBM 1311 Disk Storage Drives decrease access time for reading and writing data. These features also increase the flexibility of the program by permitting scanning, and variations of reading and writing of data in the move or load modes.

### Track-Record Feature (Models 4 and 5 only)

The track-record special feature provides for reading or writing an entire disk track with or without the track address. A single, 6-digit address is used, followed by 2,980 characters in the move mode and 2,682 characters in the load mode. Track records can be used for storing programs, tables, blocked records, and other data requiring a single large storage block.

When this feature is installed on the master drive, it also functions with the Model 2 drives attached to the system.

### Read Disk-Track Record

#### Instruction Format.

Mnemonic	Op Code	A-address	B-address	d-character
SPS MU	M or L	%F2	xxx	R
	LU (word marks)			
A	RDTR			
	RDTRW (word marks)			

**Function.** This instruction causes data to be read from a disk track into core storage. The digit 2 in the A-address (%F2) specifies that a track-record operation is to be performed. Data is read from the disk track (2,980 characters in move mode or 2,682 characters in load mode). The additional characters read are accounted for by using the normal gap between disk sectors and the sector address positions. Reading from the disk is stopped by a group-mark with a word-mark in core storage.

Reading from the track begins following the address specified by the core-sector address. This address is located at the beginning of the track, directly after the index pulse.

The core-sector address field in core storage is not modified, and the sector-count field in core storage is reduced by one as the track is read. The sector-count field must be set at 001 before the operation begins so that reducing it by one can signal an end-of-operation (000 in sector-count field).

The B-address specifies the high-order position in core storage of the disk-control field, and the area in storage reserved for data read from the disk track.

The R in the d-character position signifies that this is a read operation.

**Word Marks.** A group-mark with a word-mark must be one position to the right of the last position reserved in core storage for the track record. If a group-mark with a word-mark is detected before reading of the track is completed, the wrong-length record and any disk-condition indicators turn ON and reading stops. The position of the group-mark with a word-mark is determined by adding the B-address to 2991.

**Timing.**  $T = N(L_T + 1) + 40 \text{ ms} + \text{disk rotation}$ .  
 $N = .0115 (1401), .006 (1460), .0111 (1440)$ .

**Note.** Track-record read operations can be performed only on a track written with a track-sector instruction.

Before reading starts, an automatic check is made of the record address in storage with the record address on the disk. If the addresses are not the same, the unequal-address compare and any disk-condition indicators are turned ON, and the data in storage cannot be read from the disk.

#### Address Registers After Operation.

I-Add. Reg.	A-Add. Reg.	B-Add. Reg.
NSI	B + 6	B + 11 + 2980
		or
		B + 11 + 2682

**Example.** Read disk track 012540 in core storage beginning at location 976 (area is labeled TRSEC1). The high-order position of the disk address is in the first ten positions of the label (966-975), Figure 11.

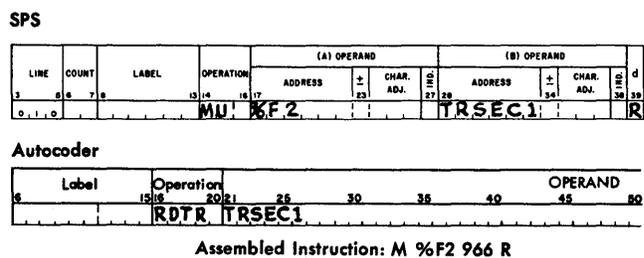


Figure 11. Read Disk-Track Record

## Write Disk-Track Record

### Instruction Format.

Mnemonic	Op Code	A-address	B-address	d-character
SPS MU	M or L	%F2	xxx	W
	LU (word marks)			
A	WDTR			
	WDTRW (word marks)			

**Function.** This instruction causes data from core storage to be written on a disk track. The digit 2 in the A-address (%F2) specifies that a track-record operation is to be performed. An entire disk track is written from the data in core storage (2,980 characters in move mode or 2,682 in load mode). The additional characters are accounted for by writing in what is normally the gap between disk sectors and the sector address positions. Writing of the disk track is stopped by sensing a group-mark with a word-mark in core storage and the end of track.

Writing begins at the track address specified by the core-sector address field. This address is located at the beginning of the track, directly after the index pulse.

The core-sector address field in storage is not modified, and the sector-count field in core storage is reduced by one as the track is written. Set the sector-count field to 001 before the operation begins so that reducing it by one can signal an end-of-operation (000 in sector-count field).

The B-address specifies the high-order position in core storage of the disk-control field, and the area in storage where the data to be written on the disk track is stored.

The W in the d-character position signifies that this is a write operation.

**Word Marks.** A group-mark with a word-mark must be one position to the right of the last character of the data in core storage. The writing of data stops when the end of track is reached on the disk and a group-mark with a word-mark is sensed in core storage. If the group-mark with a word-mark is sensed before the end of track, the remainder of the disk track is filled with data from core storage, and the wrong-length record and any disk-condition indicators are turned on.

**Timing.**  $T = N (L_1 + 1) + 40 \text{ ms} + \text{disk rotation}$ .

**Note.** Before writing starts, an automatic check is made of the core-sector address in storage with the record address on the disk. If the addresses are not the same, the unequal-address compare and any disk-condition indicators are turned on, and the data in storage cannot be written on the disk.

A WRITE DISK CHECK instruction must be performed following a write operation unless an error occurred during the write operation. No other disk-storage operation can be performed until the check of data written on the disk is accomplished.

If the data in core storage contains characters with word marks and the write operation is performed in the move mode, only the CBA 8421 portion of the character is written on the disk (the word mark is ignored).

Disk tracks adjacent to, but not above or below, a disk track written with the WRITE DISK-TRACK RECORD instruction must be either unused or set up as a *track record*. If the adjacent tracks are written using WRITE DISK SECTOR or WRITE DISK SECTOR WITH ADDRESSES instructions, interference occurs to the track-record data stored in what is normally the gap between sectors.

### Address Registers After Operation.

I-Add. Reg.	A-Add. Reg.	B-Add. Reg.
NSI	B + 6	B + 11 + 2980
		or
		B + 11 + 2682

**Example.** Write a disk track record from the data in the core-storage area labeled TRSEC1 (the first position of data is at 976). The high-order position of the disk address is in the first ten positions of the label (966-975) (Figure 12).

SPS

LINE	COUNT	LABEL	OPERATION	(A) OPERAND			(B) OPERAND			d
				ADDRESS	CHAR. ADJ.	CHAR. ADJ.	ADDRESS	CHAR. ADJ.	CHAR. ADJ.	
0	1	0	MU	%F2			TRSEC1			W

Autocoder

Label	Operation	OPERAND
WDTRW	TRSEC1	

Assembled Instruction: M %F2 966 W

Figure 12. Write Disk-Track Record

## Read Disk-Track Record with Address

### Instruction Format.

Mnemonic	Op Code	A-address	B-address	d-character
SPS MU	M or L	%F@	xxx	R
	LU (word marks)			
A	RDTA			
	RDTAW (word marks)			

**Function.** This is similar to the READ DISK-TRACK RECORD instruction except that the @ in the A-address (%F@) specifies that the address of the track record in disk storage is to be read into core storage with the data on the disk track. Data is read from the disk track (2,986 characters in move mode or 2,688 characters in load mode). The additional characters read are accounted for by using the normal gap between

disk sectors and the sector address positions. Reading from the disk is stopped by a group-mark with a word-mark in core storage.

When a disk-track-record operation is initiated, an automatic check is made of the record address in storage with the record address on the disk. If equal, reading begins immediately following the index pulse on the disk track, which signals the system that the beginning of a track is about to come under the access assembly. The track record address in the high-order position of the disk-data field in core storage is written in the first sector address position after the index pulse.

The core-sector address field in core storage is not modified. The sector-count field in core storage is reduced by one as the track is read. The sector-count field must be set at 001 before the operation begins so that reducing it by one can signal an end-of-operation (000 in sector count field).

The B-address specifies the high-order position in core storage of the disk-control field, and the area in storage reserved for the track record address and data read from the disk track.

The R in the d-character position signifies that this is a read operation.

**Word Marks.** A group-mark with a word-mark must be one position to the right of the last position reserved in core storage for the disk track. If a group-mark with a word-mark is detected before reading of the track is completed, the wrong-length record and any disk-condition indicators turn ON and reading stops. The position of the group-mark with a word-mark is determined by adding the B-address to 2997.

**Timing.**  $T = N(L_T + 1) + 40 \text{ ms} + \text{disk rotation}$ .

**Note.** Track-record read operations can be performed only on a track written with a track-record instruction.

Before reading starts, an automatic check is made of the record address in storage with the record address on the disk. If the addresses are not the same, the unequal-address compare and any disk-condition indicators are turned on, and the data in storage cannot be read from the disk.

**Address Registers After Operation.**

<i>I-Add. Reg.</i>	<i>A-Add. Reg.</i>	<i>B-Add. Reg.</i>
NSI	B + 9	B + 11 + 2986
		or
		B + 11 + 2688

**Example.** Read the address and data from disk track 012540 into core storage beginning at location 476 (area is labeled TRECAD). The high-order position

of the disk-control field is in the first ten positions of the label (466-475) (Figure 13).

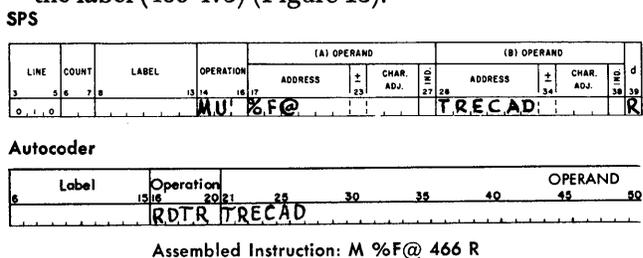


Figure 13. Read Disk-Track Record with Addresses

**Write Disk-Track Record with Address**

**Instruction Format.**

<i>Mnemonic</i>	<i>Op Code</i>	<i>A-address</i>	<i>B-address</i>	<i>d-character</i>
SPS	MU	M or L %F@	xxx	W
		LU (word marks)		
A	WDTA			
		WDTAW (word marks)		

**Function.** This is similar to the WRITE DISK TRACK RECORD instruction except that the @ in the A-address (%F@) specifies that the address of the track record in core storage is written on the disk.

When a disk-track-record operation is initiated, an automatic check is made of the record address in storage with the record address on the disk. If equal, writing begins immediately following the index pulse on the disk track, which signals the system that the beginning of a track is about to come under the access assembly. The track-record address in the high-order position of the disk-data field in core storage is written in the first sector address position after the index pulse.

An entire disk track is written from the data in core storage (2,986 characters in move mode or 2,688 in load mode). The additional characters are accounted for by writing in what is normally the gap between disk sectors and the sector-address positions. Writing of the disk track is stopped by sensing a group-mark with a word-mark in core storage and the end of track.

The core-sector-address field in storage is not modified, and the sector-count field in core storage is reduced by one as the track is written. The sector-count field should be set to 001 before the operation begins so that reducing it by one can signal an end-of-operation (000 in sector-count field).

The B-address specifies the high-order position in core storage of the disk-control field. It also specifies the area in storage where the address and data to be written on the disk track are stored.

The W in the d-character position signifies that this is a write operation.



field that specifies the starting address in disk storage. The record area associated with the disk control field contains the search argument. The search argument must be placed in the same positions of the core-storage record as it appears in the disk-storage record. Skip codes (\$) are used in those positions of the core-storage record that are not a part of the search argument (Figure 15). The search argument can be variable in length, but must be no longer than 99 characters. The units positions of the search argument should be followed by a group-mark with a word-mark to signal the end of the search argument.

The sector-count field is set to the number of sectors the programmer thinks it may require to find the desired disk record. The sector-count field is modified by minus one before each sector is scanned. The core-sector address field is modified by plus one after each sector is scanned. It is modified only if

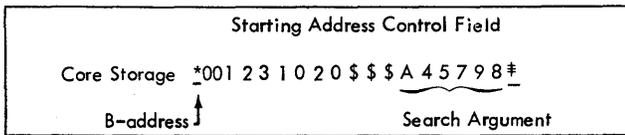


Figure 15. Record in Core Storage for Scan Operation

the sector does not satisfy the search argument. Therefore, when scanning stops, the core-sector address contains the address of the desired disk record.

Scanning begins at the disk record specified by the B-address and ends:

1. When the specified comparison is found. The sector-count field may, or may not, be all zeros at this time.
2. When the operation reaches the end of a cylinder. The sector-count field may, or may not, be all zeros at this time.
3. When the sector-count field is reduced to all zeros.

**Word Marks.** A group-mark with a word-mark must be set one position to the right of the last character of the search argument.

**Timing.**  $T = N (L_I + 1) + 2N_S + \text{disk rotation}$ .

400 ms is the maximum time for scanning one cylinder (200 sectors).

**Notes.** The High-Low-Equal Compare feature is a prerequisite for this instruction.

The result of the scan is determined by testing the high, low, or equal compare indicators with the BRANCH IF INDICATOR ON instruction.

The scan operation can be performed only on disk records written in sector format.

**Address Registers After Operation.**

I-Add. Reg.	A-Add. Reg.	B-Add. Reg.
NSI	B + 6	B + 11 + L <sub>F</sub>

**Example.** Scan disk storage for an equal compare beginning at sector-address 012510 and continue scanning until the record with part number A24537 is found. The disk-control field is located in the high-order positions of the area of core storage labeled SCANAR (966-974), Figure 16.

SPS

LINE	COUNT	LABEL	OPERATION	(A) OPERAND				(B) OPERAND				d	
				ADDRESS	Σ	CHAR. ADJ.	Q	ADDRESS	Σ	CHAR. ADJ.	Q		
3	0	0	MU	%F8				SCANAR					W

Autocoder

Label	Operation	OPERAND
SDE	SCANAR	

Assembled Instruction: M %F8 966 W

Figure 16. Scan Disk Equal

**Direct Seek Feature (Model 4 only, Standard on Model 5)**

This special feature reduces access time (250 ms maximum, 150 ms average) by allowing the access assembly to be positioned directly at a new setting without returning to the home position.

The instruction used for direct seek operation is the same as that used with normal seek (M %F0 BBB R). The B-address position of the instruction contains the core-storage address of the high-order position of the six-digit disk control field used.

**DISK-CONTROL FIELD**

Direct Seek operations use a six-position disk-control field (Figure 17).

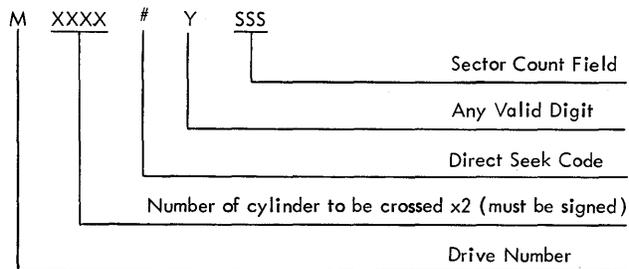


Figure 17. Disk Address for Direct Seek

The first position of the disk-control field contains the disk-drive number (0, 2, 4, 6, or 8). An asterisk cannot be used for this operation.

The next four positions (2-5) contain a signed four-digit number equal to twice the number of cylinders to be advanced (+) or retracted (-).

The sixth position contains a pound (#) sign to indicate a direct seek operation.

**Note:** Any character with 8-2-1 bit combination will be taken to indicate a direct seek operation (Figure 18).

DIRECT SEEK (8 2 1) #

CHARACTER	BIT CODE	
.	BA 8 21	} May not be used in X4 position or will get a Direct Seek
≠	BA 8421	
\$	B 8 21	
△	B 8421	
,	A 8 21	
≡	A 8421	
√	8421	

Figure 18. 8 2 1 Bit Combinations

The signed difference field can be calculated by the method shown in Figure 19. This method uses the four high-order positions of the disk address at which the access arm is positioned and the four high-order positions of the disk address to be sought. Both fields must be changed to either odd or even (add one to even, subtract one from odd). The old address is then subtracted from the new address. The result of the subtraction has the correct sign to indicate that the mechanism is to advance (+) or retract (-).

New Address	004372	0043	both odd	0043
			+ 1	
Old Address	003291	0032		<u>0033</u>
			difference	0010
Signed Difference	0010	+	(+ because increase in address)	

Figure 19. Calculated Sign Difference

DIRECT SEEK (8 2 1) #

If for some reason the address fails to specify the module in the Alternate Code position, Direct Seek will be executed on the master file.

**Direct-Seek Timing**

Figure 20 provides seek times when the Direct Seek feature is installed on the system.

**Disk-Storage Timing**

The organization of data in disk storage and the method of processing data affect the seek time for a given operation and affect also the total systems' throughput. Some methods of seeking records and the sequence of disk storage and input/output instructions are considered here as an aid to program development.

**Seeking Disk-Storage Records**

Two modes of operation for seek instructions are *Return to Home* and *Direct Seek*, a special feature. The

Number of Cylinders Traveled	Time In Milliseconds
1	54 Minimum
2	69
3	80
4	90
5	105
6	125
7	130
8	140
9	155
10	165
20	130
30	137
40	154
50	170
60	185
70	202
80	217
90	235
99	248 Maximum

Figure 20. Cylinder Seek Time with Direct Seek

Return to Home mode is the standard mode of operation. In this mode all seeks are achieved by first moving the access arms to a *home* position outside cylinder 00 and then counting into the desired cylinder. This function is automatically performed by the system. The special feature Direct Seek enables the programmer to write the program so that the system can move from one track to another track without requiring the access arms to return to a home position.

Another factor to be considered in systems planning is that the access arms move at both a low speed and a high speed. Access-arm movement within ten cylinders is at low-speed rate of two inches per second. If more than ten cylinders are used, the access arms move at the high-speed rate of 16 inches per second for all cylinders in excess of ten. These two speeds (two inches and 16 inches per second) are not used by the programmer in timing disk-storage operations because the timing charts incorporate these variations in speed. The question of variation in speed is raised here so it can be considered when data is being organized in disk storage.

After a seek instruction in either mode has been issued, processing may continue until another disk-storage instruction is issued. The length of the seek will depend on the total number of cylinders that must be passed during the seek operation. Figure 21 provides actual seek times for cylinder-to-cylinder movement in increments of ten cylinders. In the Return to Home Mode, the total throughput time can be reduced by using a technique known as *Dummy Seek to Cylinder 00*.

TO	FROM										
	00	09	19	29	39	49	59	69	79	89	99
00	75	88	101	114	127	140	153	167	179	192	204
09	175	188	201	214	227	240	253	267	279	292	304
19	143	156	169	182	195	208	221	235	247	260	272
29	153	166	179	192	205	218	231	245	257	270	282
39	168	181	194	207	220	233	246	260	272	285	297
49	184	197	210	223	236	249	262	276	288	301	313
59	200	213	226	239	252	265	278	292	304	317	329
69	215	228	241	254	267	280	293	307	319	332	344
79	232	245	258	271	284	297	310	324	336	349	361
89	248	261	274	287	300	313	326	340	352	365	377
99	263	276	289	302	315	328	345	355	367	380	392

Figure 21. Cylinder Seek Time without Direct Seek

**DUMMY SEEK TO CYLINDER 00**

Access motion time is composed of two operations: Return to Home and Advance from Home. The return-to-home portion of access time can normally be overlapped if a seek to cylinder 00 is issued prior to a card read or punch instruction or a print instruction. (See Figure 22.)

**Timing Considerations for Reading and Writing**

When designing a program utilizing the disk pack, the programmer should consider ways to place READ, WRITE, and WRITE CHECK instructions to save job time. Because the disks revolve at 1500 rpm, 40 ms are required to complete a revolution, and 2 ms to read or write one sector. The rotational time that must elapse before a disk instruction can be executed should be utilized for processing, if possible.

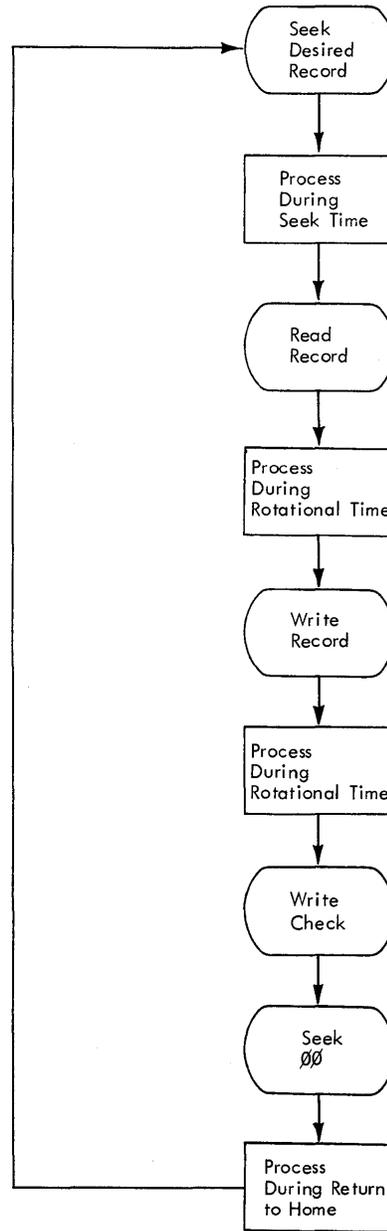
Assume, for example, that a two-sector record (200 characters) is to be read, updated, and then returned to the file. The timing chart and block diagram for this operation are shown in Figure 23.

The total time for this operation is 106 ms, for approximately 2 5/8 disk revolutions. The available processing time is 68 ms.

Processing time is reduced as more sectors are read or written. The timing for a 4-sector operation illustrates this point:

- Read            2 ms head select delay time  
                  20 ms average rotational time  
                  8 ms to read 4 sectors
- Process        30 ms processing

- Write            2 ms head select delay time  
                  8 ms to write 4 sectors
- Process        30 ms processing
- Write Check    2 ms head select delay time  
                  8 ms write check  
110 ms total



Reading, punching, or printing may be done during the pertinent record seeking and the seek 00 time.

Figure 22. Block Diagram for Dummy Seek Technique

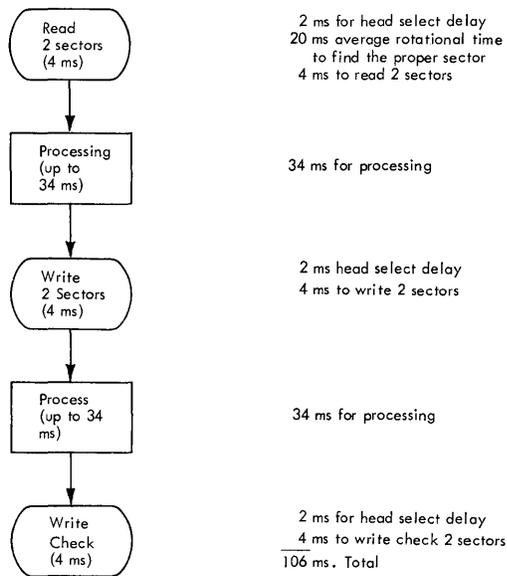


Figure 23. Disk Storage Timing for a Two-Sector Record

The total time in the preceding example is 110 ms ( $2\frac{3}{4}$  revolutions), only 4 milliseconds longer than the 2-sector operation. However, total processing time is 60 ms as opposed to 68 ms in the earlier example.

If possible, processing should be kept within the available rotational time. If not, the cycle will be in-

creased by one 40-ms revolution for each extension of available processing time.

Processing time between a WRITE instruction and a WRITE CHECK instruction can be used for updating control totals and/or arranging fields of printing. When the Print Storage special feature is installed, most disk operations may be completely overlapped by the printing operation.

A summary of the disk-storage times follows:

Rotational Delay	40 ms
Average Rotational Delay	20 ms
Head Select Delay	2 ms
Read One Sector	2 ms
Write One Sector	2 ms
Write Check One Sector	2 ms

#### Seek Time—without Direct Access

Maximum	400 ms
Mean Seek Time	250 ms

#### Seek Time—with Direct Access

Maximum	250 ms
Mean Seek Time	150 ms



**International Business Machines Corporation  
Data Processing Division  
112 East Post Road, White Plains, N. Y. 10601**