

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, is positioned within a black rectangular box.

Systems Reference Library

Disk Utility Programs Specifications IBM 1401, 1440, 1460 (with IBM 1301 and 1311)

Program 1401-UT-053
Program 1440-UT-041

This publication contains the specifications of nine disk-storage utility programs for IBM 1401, 1440, and 1460 Systems equipped with either or both 1301 and 1311 disk storage. The programs are:

- Clear Disk Storage
- Disk to Tape
- Tape to Disk
- Disk to Card
- Card to Disk
- Copy Disk
- Print Disk
- Disk Record Load
- Disk Label

For a list of related publications and abstracts, see the publications IBM 1401 and 1460 Bibliography, Form A24-1495, and IBM 1440 Bibliography, Form A24-3005.

Fourth Edition

This is a reprint of C24-1484-2 incorporating changes released in the following Technical Newsletters:

<u>Form No.</u>	<u>Pages Affected</u>	<u>Date</u>
N24-0287	7, 8	11/25/64
N21-0054	5, 7, 15, 17, 18, 28	9/1/65
N21-5001	14, 17, 18, 28, 29	12/12/65

Requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form is provided at the back of this publication for reader's comments. If the form has been removed, comments may be addressed to IBM Corporation, Programming Publications, Dept. 425, Rochester, Minn. 55901.

CONTENTS

INTRODUCTION	5
General Description	6
Control Cards	6
Area-Control-Card Check	10
Error Procedures	10
 CLEAR-DISK-STORAGE PROGRAM	 12
 DISK-TO-TAPE PROGRAM	 14
 TAPE-TO-DISK PROGRAM	 17
 DISK-TO-CARD PROGRAM	 19
 CARD-TO-DISK PROGRAM	 23
 COPY-DISK PROGRAM	 24
 PRINT-DISK PROGRAM	 26
 DISK-RECORD-LOAD PROGRAM	 28
 DISK-LABEL PROGRAM	 29
Specific Operations	29
 TIMINGS	 32

DISK UTILITY PROGRAMS SPECIFICATIONS
IBM 1401, 1440, 1460 (WITH 1301 AND 1311)

Whatever may be the specific uses to which a data processing system is put, there exist certain unique operations that must be performed frequently. These operations may differ in detail, depending on the particular machine configuration and data format of the individual user, while the essential function remains the same. Because of their frequent use, the burden of programming these operations for each specific, and perhaps non-recurring, job could be prohibitive. Therefore, there is a need for generalized routines designed to satisfy specific functions. These routines must be flexible enough to allow the user to assign the specifications of his particular problem.

IBM supplies several types of programs that meet these requirements. Those described in this publication are grouped under the heading, Disk Utility Programs. They are designed to assist the user in the day-to-day operation of his installation. By means of these programs, certain frequently required operations, such as loading or unloading disk files from cards or tape, and printing out areas of disk storage for program testing purposes, can be performed without programming effort on the part of the user.

This publication describes two packages of disk utility programs supplied by IBM. One of the packages is programmed to operate on 1401 and 1460 systems equipped with either 1301 or 1311 disk storage. The other package is for 1440 systems equipped with either 1301 or 1311 disk storage. The specifications of the programs are the same for both packages, with very few exceptions.

Each package includes the following programs:

Clear-Disk-Storage Program
Disk-to-Tape Program
Tape-to-Disk Program
Disk-to-Card Program
Card-to-Disk Program
Copy-Disk Program
Print-Disk Program
Disk-Record-Load Program
Disk-Label Program

The clear-disk-storage program clears an entire disk pack or module, or any specified portions of it, by filling these areas with blanks or any other valid 1400-series character.

The disk-to-tape program writes the contents of specified areas of disk storage on magnetic tape. This data can be reloaded by the tape-to-disk program.

The disk-to-card program punches the contents of specified areas of disk storage into cards. This

data can be reloaded by means of the card-to-disk program.

The copy-disk program writes the contents of specified areas of disk storage into other areas on the same unit or on a different unit.

The print-disk program enables the user to print out the contents of any specified areas of disk storage.

The disk-record-load program enables the user to insert any desired number of characters into any location in disk storage. The data can be inserted in either the address or record portion of a sector or track record.

The disk-label program is not used for 1301 disk storage. It is used on a 1311 to set up the initial header labels on a disk pack, and to perform necessary maintenance operations on existing labels.

These programs may be loaded back-to-back, in any combination, for continuous operation.

Minimum Machine Requirements

1401-1460 Disk Utility Programs

- 4,000 positions of core storage (8,000 positions required if the track-record feature is used)
- One IBM 1402 Card Read-Punch
- One IBM 1403 Printer, Model 1, 2, or 3
- One IBM 7330 or 729 Magnetic Tape Unit (tape programs only)
- Either:
 - One IBM 1311 Disk Storage Drive, or
 - One IBM 1301 Disk Storage, Model 11

1440 Disk Utility Programs

- IBM 1441 Processing Unit with 4,000 positions of core storage (8,000 positions required if track-record format is used)
- One IBM 1443 Printer
- One IBM 7335 Magnetic Tape Unit (tape programs only)
- Either:
 - One IBM 1442 Card Read-Punch, Model 1 or 2, or
 - One IBM 1442 Card Reader, Model 4, and one IBM 1444 Card Punch, Model 1.

- Either:
One IBM 1311 Disk Storage Drive, or
One IBM 1301 Disk Storage, Model 11.

GENERAL DESCRIPTION

These nine utility programs are contained in self-loading, condensed card decks. No Autocoder or symbolic assembly is necessary for their operation. The disk-label program is contained within one card deck. Each of the other eight program decks is made up of four sections (Figure 1).

Loader. The first section of each program deck is a self-loading program called the loader. The 1401/1460 programs use a six-card loader. The 1440 programs use a seven-card loader. The loader clears core storage and loads the rest of the program deck. This loader performs the same function, but is not the same as, the loaders generated for other programs assembled by 1401 Autocoder.

Disk-Header-Label Routine. This routine applies to 1311 disk storage only. Each of the eight programs includes a routine that performs 1311 disk-label operations appropriate to the particular application. For example, the disk-header-label routine included with the clear-disk-storage program checks to see that there are no unexpired files within the area specified. If all of those labels are expired, they are printed and then deleted from the label track.

In certain cases, the disk-header-label routines can be used to generate the information normally required in the area-control cards.

The disk-header-label routines perform their operations in accordance with specifications punched in one or more header-control (RDLIN) cards. The routine is loaded and executed before the main program is loaded. If no operations are to be performed on the label track during a particular program run, no control cards need be supplied to the routine. When the user chooses not to implement disk-label procedures in a 1311 installation, or when he is using 1301 disk storage, the header-label routines can be removed from the program decks to allow faster program loading.

4K Object Program. This is the basic program, complete with end card. This section with the loader is sufficient for operation of the program. A user with no more than 4,000 core-storage positions must remove the cards following the 4K end card before loading the program.

8K Modification to 4K Object Program. This section is used in a system with 8,000 or more core-storage positions. It includes the instructions necessary to:

1. operate in the track-record mode,
2. use alternate tape drives for multi-tape files, and
3. check tape-header labels on other than the first file of multi-file tapes.

If none of these operations is performed, this section need not be used. If this section is used, remove the 4K end card from the object deck.

CONTROL CARDS

Three types of control cards are used with these programs. The first are RDLIN (Read Label Information) cards used to provide information to the disk-label program and to the disk-header-label routines of the other programs. Second is the system-control card that is used to bypass halts, select the number of tape drives, choose the type printer to be used, determine if the first 80 or all 120 characters of the standard 120-character (type A) tape label are to be checked, and for 1440 systems, select the punch-unit numbers. Third are area-control cards that provide information to the main program. The information contained in these area-control cards is used to modify the program to fit the particular specifications for each processing run. For instance, the operation codes of disk operations must be established as either M or L, depending on whether the operations are to be in the move or load mode. The disk-control fields must

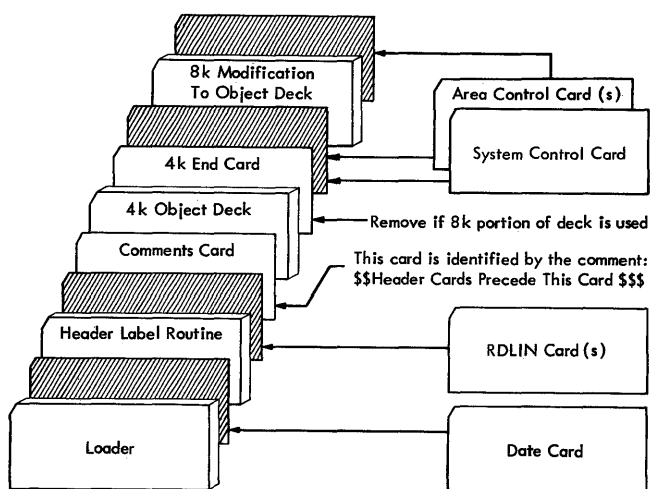


Figure 1. Program-Deck Makeup

be set up, and must specify the disk drive to be used and the sector address at which the operation is to begin. Linkage to certain optional routines must be either established or broken, depending on the user's specific requirements.

In addition to these three control cards, a date card is used by the header-label routine of some programs.

Disk-Label Control Cards

The control cards used with 1311 disk-header-label routines are similar to the RDLIN information cards used with the 1401 IOCS RDLIN (Read Label Information) macro instruction. Figure 2 illustrates the format of these cards.

The first 15 columns contain control information that specifies the disk drive(s) being used and, in some cases, the operation to be performed. If the disk-header-label routine is to supply the area limits to the program, these columns also contain optional information normally contained in the area-control cards:

Column 16-20 must contain the entry RDLIN.

Columns 21-66 contain information as it appears in the disk labels to be checked. The label fields required vary with different programs. In some cases, only the area limits are used; in others, only the file identification. Each field must be punched in the columns indicated in Figure 2.

The RDLIN cards are placed after the disk-header-label routine and before the main program deck.

Date Card

If the retention cycle is to be checked by the disk-label routine or the tape label routine of a program, a date card must be supplied by the user. This card contains the date of the day that the program is being run. The card is punched as follows:

Columns	Contents
1-3	082 (Location of date field in core storage)
4-5	05 (Length of date field)
6	word separator (0-5-8 punch)
7-8	year (i.e., 63)
9-11	day of year (001 represents Jan. 1, 365 represents Dec. 31)

This card is inserted between the loader and the header-label-routine deck. The date is loaded into core-storage locations 82-86.

System-Control Card

The system-control card is used to specify the type of equipment attached to the system, the halts that are to be bypassed, and the choice of tape-label checking to be done. Normally the card need be punched only once, and can be used for all eight programs. The card is punched as follows, and any punch other than those described will be treated as a blank. Figure 3 illustrates the format of the card.

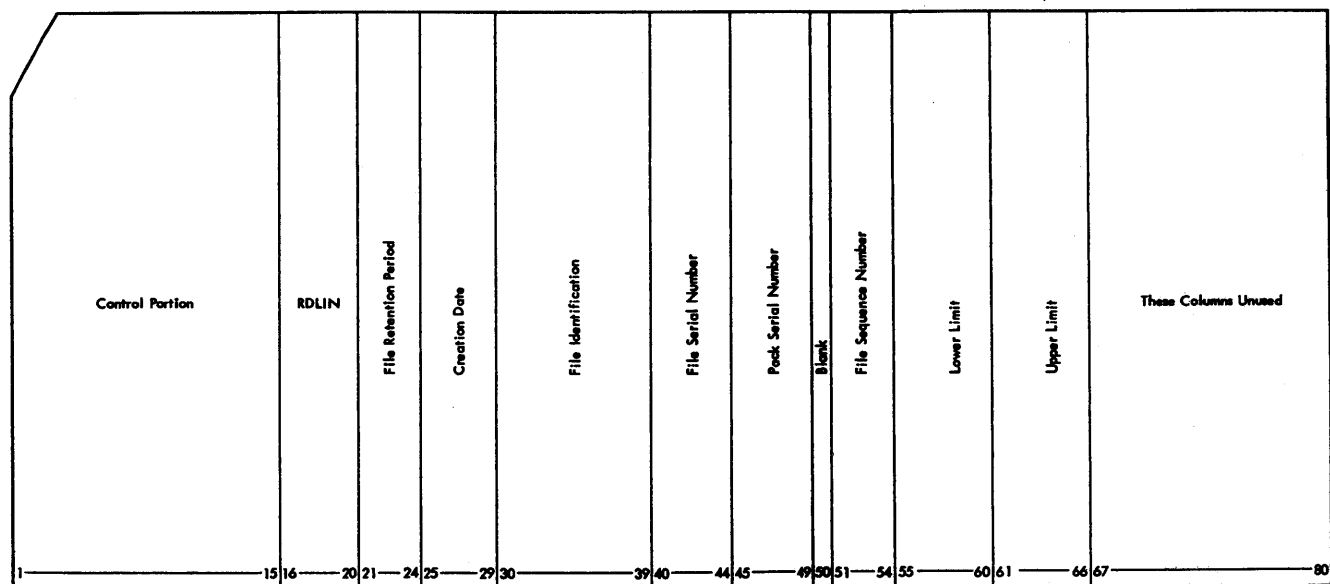


Figure 2. RDLIN Card

Columns	Contents	Reason
1-3	CTL	Card Identification
4	1	A 1-punch indicates that disk and tape read error halts are to be bypassed. A blank means halt on error.
5	1	A 1-punch indicates that area control-card analysis halts are to be bypassed. A blank means halt after analysis.
6	1 or 2	Denotes the number of tape drives to be used. (Two drives may be used only if the 8K portion of the deck is used.) A blank means one tape unit.
7	1 or 2	A 1-punch denotes a 100-position printer. A blank or 2-punch denotes a 132-position printer. 120 positions are always assumed for a 1440 system.
8	1, 2, or 3	Denotes the punch unit select number. A blank means unit 1.
9	Blank	This column is not used.
10	1 or 2	A 1-punch indicates that only the first 80 positions of the standard 120-character tape-label are to be checked. A blank or a 2-punch indicates that all 120 characters are to be checked.

Error-Halt Option

Normally, these programs halt after an uncorrectable disk or tape error. However, the user can specify by a 1-punch in column 4 of the system-control card, in the event of a disk or tape read error, that the program bypass the track or area on which the error occurred, print an error message, and continue processing without halting. All disk-error halts are bypassed except access-inoperable errors. Card-error halts cannot be bypassed.

CTL	Disk and Tape Error Halts	Control Analysis Halts	Number of Tape Units	1403 Print Positions	Punch Unit Select Number	Blank	80-120 Positions Tape Label
1-3	4	5	6	7	8	9	10

Figure 3. System-Control Card

Analysis-Halt Option

Each of these programs has similar control-card analysis routines, described under Control-Card Check. Normally, the programs halt after printing the control-card analysis. However, the user can specify by a 1-punch in column 5 of the system-control card that these halts are not to occur.

Area-Control Cards

The disk-label program and the disk-record-load program do not use separate area-control cards. The other seven programs use cards punched as illustrated in Figure 4.

The information required in the area-control cards for each of these programs includes:

1. the upper and lower limits of the disk-storage areas to be operated on.
2. the mode of operation to be used (M or L).




Figure 4. Area Control Card

3. the disk drive(s) or module(s) being used.
4. an indication of the track format used in each area.

The rules for determining the information are similar for all of the programs and are described in this section. Certain additional unique information is required for each individual program and is described in the separate program descriptions. The entire input formats required for the disk-record-load program and the disk-label program are described separately.

Under certain conditions, the 1311 disk-header label routines used with the disk-to-card, disk-to-tape, copy-disk, and print-disk programs can extract the area limits from the disk labels of up to four files to be processed. In this case, the area-control cards are not required. The disk-header-label routines place the control information for up to four areas in the normal area-control-card read-in area and then transfer control to the main program. Additional areas can be defined with area-control cards.

Area Definition

As many as four areas can be defined in one control card. If more than four areas are to be operated upon, as many additional control cards as are required can be used. A single area must consist of at least one full track, and can consist of as many tracks as desired, if:

1. they are consecutive tracks within a single disk pack or module
2. they are all to be processed in the same mode
3. the same options are used on each track in the area.

Thus, with a single area-control card, these programs can operate on as little as a single track or as much as four full disk packs or modules.

Lower Limit. The lower limit of each area specified must be the address of the zero sector of the first track to be operated upon. This is true regardless of the addressing structure of the area. For example, if the first track of an area had only five unique sector addresses, 000045, 000046, 000047, 000048, and 000049, and these were repeated four times around the track, the address given as the lower limit of the area would be 000040.

Upper Limit. The upper limit of an area normally should be the address of the last sector on the last track of the area. The following discussion of the way in which this address is used is important only in the case where unusual addressing schemes are used.

The upper limit of an area can be any sector address that falls within the sequence of addresses valid for the last track. Normally, the address of the last sector in the area should be used. However, if the addressing scheme is not consistent throughout the area, the address of a different sector might be used.

A sector address relative to the one given is used to operate on every track in the area. Suppose, for example, that the lower and upper limits of an area are given as 000040 and 000398, respectively. The program repeatedly subtracts 20 from the two low-order positions of the upper limit (98), until the result of the subtraction is less than 20. In this example the remainder would be 18. This remainder is added to the lower limit. The sum, 000058, is the address with which the program seeks and attempts to read or write the first track in the area. If an unequal-address compare results (indicating that this address was not found on the track), the program resets the address to that of the zero sector of the track, and attempts the operation again. If this address is not found, a constant is added to the address.

This process is repeated until an address is found to compare with one of those on the track or until all valid addresses for that track have been tried. The second track, in this example, would be searched first for the address, 000078. If it is not found, the process is repeated again for this track. Thus, if the addressing structure is known, the upper limit should be the address of a sector that is relative to sector addresses actually used on tracks throughout most of the area. Further, the upper limit should be an address that is located as far past the index point as possible. This minimizes rotational delay time.

However, because the address specified need not be one of those actually used on the track, it is possible to process an area even though the exact addressing structure is unknown.

Mode of Operation

Each of these programs, with the exception of the disk-label program, can operate in either the move or load mode. An M or an L must be punched for each area defined, to indicate the mode to be used for that area. An area of disk storage must be read in the mode in which it was written.

Disk-Unit Character

The character of the disk drive or module on which each area is located must be specified. These characters are:

- 0 for the first disk drive in the system
- 2 for the second drive
- 4 for the third
- 6 for the fourth
- 8 for the fifth.
- F for the first module in the system
- S for the second module
- U for the third
- W for the fourth
- Y for the fifth.

Track-Record Operations

A 1311 or 1301 disk track is normally written with 20 sectors. Each sector contains a 6-digit address and a record area of 100 characters in the move mode or 90 characters in the load mode. This is referred to as the 20-sector format.

A special feature is available for the 1400 series, making it possible to write areas using the track-record operation, where only one address and one record area are used per track. The 1311 track contains one 6-digit address and a single record area of 2,980 characters in the move mode or 2,682 characters in the load mode. The 1301 track contains one 6-digit address and a single record area of 2543 characters in the move mode or 2261 characters in the load mode. This is referred to as the track-record format.

For each area defined in the area-control cards, one column is used to indicate the track format in that area. If the area is written in 20-sector format, this column is left blank. If the area is written in the track-record format, this column contains a 1-punch.

Note: If this feature is used, 8,000 positions of core storage are required.

AREA CONTROL-CARD CHECK

After a program is loaded into core storage and the system-control card has been analyzed, it automatically reads in the first area control card. The program prints out the upper and lower limits of the area to be processed, and an analysis of the other information in the card. The program halts at this time, unless it has been indicated in the system-control card to bypass this halt. If the program

halts, the user can check the analysis printout and choose one of two options:

1. Press START to process these areas
2. Press START RESET and START to bypass these areas and read another area control card.

This process is repeated for each control card used.

ERROR PROCEDURES

Disk-Operation Errors

In the event of an error in performing a disk operation, these programs go into error routines that attempt the operation up to two more times. An unequal-address compare causes the program to attempt the operation using each of the 20 valid addresses for the track. If the error persists, the programs print the address of the track on which the error occurred and a message indicating the type of error. The program then halts, unless the user has indicated in the System-Control card that he does not wish the halts to occur. If a halt occurs, the user can press:

START, to cause the operation to be attempted three more times or

START RESET and START, to cause the program to continue as if the halt had been bypassed.

Bypassing a halt causes the following programs to skip to the next track:

Clear-Disk Storage

Tape-to-Disk

Card-to-Disk

Copy-Disk

Print-Disk

The disk-to-tape and disk-to-card programs skip to the next area specified. The disk-record-load program skips to the next card. Disk-error halts encountered by the disk-label program cannot be bypassed. The operation is not reattempted before the halt. The user can press START to try the operation again.

Tape-Write Errors

In the event of an error in attempting to write a tape record, the disk-to-tape program attempts to write the record one more time, and then it executes a SKIP AND BLANK TAPE instruction. This process is repeated up to 25 times, or until the record is successfully written. If unsuccessful after 25 tries, the program prints a message indicating that the tape is unusable. It then writes double tape marks and performs a rewind and unload operation. The user can load a new tape and continue to end-of-job.

Tape-Read Errors

In the event of an error in attempting to read a tape record, the tape-to-disk program backspaces and attempts to read the record 99 more times. If the error persists, the program prints the track address of the record. Depending on the user's choice, the program then either halts or attempts to write the erroneous record on the disk and continue. If the user chooses to have the program halt, he has the following options:

1. try the read operation 100 more times
2. correct the record in core storage
3. bypass the record
4. write the erroneous record on the disk.

Card-Read Errors

In the event of a card-read error, the machine halts. The user can replace, correcting if necessary, the card that caused the error, and press START to reread the card.

Card-Punch Errors

In the event of a card-punch error when using a 1401 or 1460 system, the disk-to-card program attempts to punch the card again. If unsuccessful after four attempts, the program halts. When using a 1440 system the program prints an error message and halts after every punch error. The last card in the stacker is the card in error.

CLEAR-DISK STORAGE PROGRAM

The clear-disk-storage program writes specified areas of disk storage with any desired valid 1400-series character. The program can be used whenever it is necessary to set certain areas to an initial condition. Normally before loading a file onto disks, the file area is written with blanks. In some cases, however, it is desirable to use some other character as a padding character. When a track or record is subsequently read from the disk, unused locations are easily identified as fields containing only the clearing character.

An area can be cleared in either the move or load mode, and in either 20-sector or track-record format. An area should be cleared in the mode and format that will be used subsequently to read from that area.

ADDRESSES

This program writes the sector addresses in the address portions of the areas cleared. The user has two options from which to choose in writing the addresses.

1. Sequential addresses can be written, beginning with the address given as the lower limit of the area, and continuing through the track containing the upper limit. If the track-record format is used, addresses are written relative to that given as the upper limit of the area. For example, if the upper limit is given as 019981, the addresses on the first five tracks would be: 000001, 000021, 000041, 000061, 000081. Neither the mode nor the format used need be that in which the area was previously written. The pack or module can have any previous addressing reference. That is, the addresses on the disk pack or module before the operation can be those normally associated with any disk drive or module.
2. The same six-digit addresses as those originally in the area can be rewritten. If this option is taken, the area must be cleared in the same mode (M or L) and the same format (20-sector or track-record) as that in which it was previously written. This operation is indicated by a dollar sign (\$) in column 15 of each area definition.

AREA-CONTROL CARDS

Each 20-column field used to define an area to be cleared is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	M or L to indicate move or load mode
2-7	Lower limit
8-13	Upper limit
14	Disk drive or module being used
15	Disk drive or module to which addresses are referenced (if retaining same addresses, punch \$)
16	The character to be used in clearing
17	Track format blank = 20-sector 1 = track record
18-20	Not used

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the clear-disk-storage program is used to insure that all files within the area to be cleared are expired, and, if so, to delete the header labels pertaining to those files.

The user punches the limits of the area(s) to be cleared in one or more RDLIN cards. The program searches the label track for labels that pertain to files within any portion of the limits provided. Labels are checked for files within the same relative limits with addresses referenced to other disk drives. For example, if the limits given were 080000-095999, the labels on the pack would be searched for files within the following limits:

080000-095999
000000-015999
020000-035999
040000-055999
060000-075999

If the retention cycle of any labels found by this search have not expired, the program prints a message to that effect along with the unexpired label(s) and then halts.

If labels pertaining to the area are expired, they are deleted by blanking the file-identification field. All labels so deleted are printed.

RDLIN CARDS

A RDLIN card must be supplied for each area to be checked. The card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	The number of the disk drive being used
2-15	Not used
16-20	RDLIN
21-54	Not used
55-60	The lower limit of the area
61-66	The upper limit of the area
67-80	Not used.

A date card must be supplied to this routine.

DISK-TO-TAPE PROGRAM

Before processing a file of disk records, it is often desirable to produce a copy of the disk file. For example, a copy might be made of a file of records before that file is updated or changed in any way. In the event of an error in processing, where the original records have been altered, the disk unit could be reloaded with the unchanged records. Having corrected the error, the processing run could then be begun again. The disk-to-tape and tape-to-disk, the disk-to-card and card-to-disk, and the copy-disk programs are used for such purposes. Each of these programs has certain advantages, depending on several factors, including the particular machine configuration available, the size of the file, the length of time the copy must be kept, etc.

The disk-to-tape program is particularly suitable when large files are to be kept for long periods of time. Also, the compact size of a tape reel makes the program especially useful when a file must be shipped from one installation to another.

However, tape is not recommended as backup storage for the autocoder systems pack, because incorrect transfer of data may result when word-separator characters are in any load-mode area on the disk.

The program unloads any number of specified areas of disk storage onto tape. An area can have been written in either the move or load mode and in either 20-sector or track-record format. Each area is written on tape in the same mode as that used in disk storage.

Normally, a single tape reel is sufficient to contain the data being unloaded in any one run. This reel will be on tape unit 1. However, if the areas to be unloaded extend through several disk packs, additional reels are required. On a 1401 or 1440 system with 4,000 positions of core storage, only one tape drive can be used. When an end-of-reel condition is sensed, the program performs a rewind and unload operation. The user can load a new tape reel and continue.

If 8,000 positions of core storage are available, a second tape drive (tape unit 2) can be used. In this case, when an end-of-reel condition is sensed, the program executes a rewind and unload operation and continues writing on the second tape drive. Successive reels are written alternately on these two tape drives.

AREA-CONTROL CARDS

If a labeled file is to be unloaded, area-control cards are not required (see Disk-Header-Label Routine). If area-control cards are used, each 20-column field used to define an area to be unloaded is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	M or L to indicate move or load mode
2-7	Lower limit
8-13	Upper limit
14	Disk drive or module being used
15,16	Not used
17	Track format blank = 20-sector 1 = track record
18-20	Not used

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the disk-to-tape program is used to insure that the proper disk file is on-line and, if so, to print and punch its header label. The label is punched in the standard RDLIN card format. This RDLIN card can be used to enter the label with the tape-to-disk program. The user must punch the drive number in column 1 before reloading.

If desired, the routine will extract the area limits from the disk label and supply these to the main program, thus making the area-control card unnecessary. This operation can be performed for up to four RDLIN cards. Additional RDLIN cards can be used for label checking only.

In addition to the area definitions supplied by the header-label routine, the program deck can be followed by as many area-control cards as desired.

The RDLIN card is punched as follows if limits are to be extracted from the label:

<u>Columns</u>	<u>Contents</u>
1	The number of the disk drive being used
2	blank
3	\$ (indicates that limits are to be gotten from label)
4	M or L to indicate move or load mode
5,6	Not used
7	Track format blank = 20-sector 1 = track record
8-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification
40-80	Not used

If the disk-header labels are to be only checked, printed and punched, with area limits supplied in area-control cards, the routine operates in one of three ways, depending on the contents of columns 21-66 of the RDLIN card.

1. If the RDLIN card contains the file-sequence number in columns 51-54, the program compares the following fields in the RDLIN card to the corresponding fields in the header labels on the track:

<u>Columns</u>	<u>Contents</u>
21-24	File Retention Period
25-29	Creation Date
30-39	File Identification
40-44	File Serial Number
45-49	Pack Serial Number
51-54	File Sequence Number

If the columns 55-66 are punched, these fields are compared to the limits in the label.

2. If the RDLIN card contains no file-sequence number in columns 51-54, but does contain the file identification in 30-39, the file identification is compared to the corresponding field of the header labels. Again, if the limits are punched, they too are compared.
3. If the RDLIN card contains no file-sequence number and no file-identification field, the limits must be punched. The header labels of all files within the limits are printed.

In addition to one of the above combinations, the RDLIN card must contain the drive number in column 1 and the entry RDLIN in columns 16-20.

The date card is not required with this particular routine.

TAPE-HEADER LABELS

If desired, this program checks and writes standard 120-character tape header labels (refer to the IBM SRL publication, Input/Output Control System on Disk for IBM 1401/1460, C24-1489. If this option is taken, a tape-header-label card with the following information must be the first card following the program deck.

Note: If 120-character tape-header labels are used, with appropriate information for 120 positions, two tape-header-label cards must follow the program deck. Information for positions 81-120 must be punched in card columns 1-40 of the second card.

120-Character Header Label

(First Card)

<u>Positions</u>	<u>Contents</u>	<u>Meaning</u>
1-5	1HDRb	Identification
7-10	DDDD	Retention cycle
11-15	YYDDD	Creation date
16-25	10 characters	File name
26-30	5 numbers	File serial number
31-35	5 numbers	Tape serial number
37-40	XXXX	Reel sequence number
41-120	Other IBM standard label information	

(Second Card)

1-40	IBM standard tape label information for positions 81-120
------	--

The program checks the first tape record on the output tape to determine:

1. Is this a standard 120-character tape-header label?
2. Is the retention cycle past?

If all these conditions are met, the program writes a new header label, using the contents of the tape-header-label card(s).

If all conditions are not met, the program prints an indication of which test(s) failed, prints the tape label received, and the header card that the user provided. The user then has the option of changing the tape or accepting the header label received. If the label is accepted, a new tape-header label is written with the information provided in the card. If a standard tape label was found, the program retains the same tape serial number that was found on the tape. If not, the serial number in the card is written.

In a system with 4,000 core-storage positions, tape header labels are checked and written only on the first tape of each run. If the user desires to process header labels on all the tapes, he must define his areas in such a way that a single area does not carry over from one reel to another. He can then reload the program for each tape reel, supplying a header card for each reel.

If 8,000 core-storage positions are available, header labels are checked and written on all output tapes if the user elects to do so on the first reel.

The tape serial number is extracted from the old label, and the reel sequence number is increased by one. The new header label is then written.

This routine must use a date card, inserted between the loader and the header-label routine deck.

TAPE-TRAILER LABELS

The program writes a standard 120-character trailer label at the end-of-job. The first 5 positions of the label contain 1EOFB, to identify end-of-file. Positions 67-72 for 120-character labels contain a block count. This count is not checked by this program or by the tape-to-disk program. It can be checked in other programs that use the tapes produced by this program as input. Positions 11-80 are left blank by the disk-to-tape program.

In a 4K system, end-of-reel is marked by double tape marks.

In an 8K or larger system, end-of-reel is marked by:

1. double tape marks, if header labels are not used.
2. a standard EOR trailer label, if header labels are used.

OUTPUT FORMAT

The first tape record produced by the disk-to-tape program (after the header label, if that option is

taken) contains the contents of the first control card. Successive tape records contain the data from the areas specified in that card. Each of these tape records contains the data from one complete disk track. The first six digits contain the address of the track being unloaded. This is the address of the zero sector of the track, regardless of the actual addresses on the track. The next three digits are the sector-count field to be used in rewriting the record on the disk by the tape-to-disk program. In the 20-sector format, this field is 020; in the track-record format, 001. This field is followed by the addresses and data from the track.

A single tape mark is written after each area.

When all the tracks specified by the first area-control card have been written on tape, the contents of the second control card (if present) are written, followed by the data from those areas. This process is repeated for each control card.

TAPE-TO-DISK PROGRAM

The tape-to-disk program reloads into disk storage the tape records produced by the disk-to-tape program described previously. The data can be reloaded into a disk pack that has any previous addressing structure, referenced to any disk drive. It is reloaded into the same relative location as that from which it was read.

The program can reload all the data unloaded by the disk-to-tape program, or it can reload selected portions of that data. The areas to be reloaded must be specified in the same sequence as that used to write the tape. If selected areas are reloaded, the program bypasses those not specified. The program can select certain areas from those on the tape, or portions of those areas. The smallest block of data that can be reloaded is that from one full track.

MULTIPLE-REEL FILES

On a 1401 or 1440 system with 4,000 positions of core storage, only one tape drive can be used (tape unit 1). When an end-of-reel condition is sensed, the program performs a rewind and unload operation, and then halts. The user can load the next tape reel and continue.

If 8,000 positions of core storage are available, a second tape drive (tape unit 2) can be used. In this case, when an end-of-reel condition is sensed, the program executes a rewind and unload operation and continues reading on the second tape drive. Successive reels are read alternately on these two tape drives.

AREA-CONTROL CARDS

Each 20-column field used to define an area to be reloaded is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Mode - M or L
2-7	Lower limit
8-13	Upper limit
14	Disk drive or module being used
15	Disk drive or module to which addresses are referenced
	<u>Note:</u> This number can be prepunched in the control cards used with the disk-to-tape program. That program does not check this column.
16	Not used
17	Track format
	blank = 20-sector
	1 = track record
18-20	Not used

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the tape-to-disk program is used to insure that all files within the disk area being loaded are expired and, if so, to enter a disk-header label for the tape file.

The user punches the limits of the area(s) to be loaded in one or more RDLIN cards. The program searches the label track for labels that pertain to any portion of the limits provided. Labels are checked for files within the same relative limits with addresses referenced to other disk drives. For example, if the limits given were 040040-055019, the labels on the pack would be searched for files within the following limits:

040040-055019
060040-075019
080040-095019
000040-015019
020040-035019

If the retention cycle of any label found by this search has not expired, the program prints a message to that effect along with the unexpired label(s), and then halts.

If all labels pertaining to the area are expired, they are deleted by blanking the file-identification field. All labels so deleted are printed out. The routine then enters the new label, as punched in the RDLIN card, on the track.

A date card must be supplied for this routine.

RDLIN CARDS

A RDLIN card must be supplied for each area to be loaded. The cards punched by the disk-to-tape program can be used. Each card must contain the limits of the area to be checked for active header labels. The disk-header-label written will contain these limits and any other fields punched in the

RDLIN card with the exception of the pack serial number. The card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	The number of the disk drive being used
2-15	Not used
16-20	RDLIN
21-29	Optional. If punched, these fields are written in the header label.
30-39	File identification
40-44	Optional. If punched, this field is written in the header label.
45-49	Not used (old pack serial number is retained).
50	Blank.
51-54	Optional. If punched, this field is written in the header label.
55-60	The lower limit of the area
61-66	The upper limit of the area
67-80	Not used

A date card must be supplied to this routine.

TAPE LABELS

If tape header and trailer labels were written by the disk-to-tape program, a tape-header card must be

used with this program also. The same card that was used in creating the tape can be used. The entire card is compared with the tape label. If there are any inequalities, the program prints both the card image and the tape label. The user can either change the reel and press START to try the comparison against the new label, or accept the label as read by turning sense switch C on and pressing START RESET and then START.

Only the header label of the first reel of multiple-reel files is checked on a system with 4,000 core-storage positions. If 8,000 core-storage positions are available, all tape-header labels can be checked.

Trailer labels on tape are used only to indicate end-of-reel. The block count is not checked. If an EOF trailer label is read before completing the areas defined in the control card, an error message is printed and the program halts.

The disk-to-card program punches the contents of specified areas of disk storage into cards. It is normally used when small areas of disk storage must be unloaded. The program can unload areas written in either the move or load mode and in either 20-sector or track-record format.

AREA-CONTROL CARDS

If a labeled file is to be unloaded, area-control cards are not required (see Disk-Header-Label Routine). If area-control cards are used, each 20-column field used to define an area to be unloaded is punched as follows:

Columns	Contents
1	M or L to indicate move or load mode
2-7	Lower limit
8-13	Upper limit
14	Disk drive of module being used
15,16	Not used
17	Track format blank = 20-sector 1 = track record
18-20	Not used

If a disk-to-card program is being punched on a second 1442, a 1444, or a 1402, the program will punch an area-control card prior to punching the data cards. This control card is a duplicate of the disk-to-card control card. The control card can be used in the card-to-disk program to load data back into the same location of disk storage it was taken from. If reading and punching is being done from the same 1442, the user's disk-to-card control card will precede the punched data cards into the stacker.

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the disk-to-card program is used to insure that the proper disk file is on-line and, if so, to print and punch its header label. The label is punched in the standard RDLIN card format. The card can be used to enter the label with the card-to-disk program. The user must punch the drive number in column 1 before reloading.

If desired, the routine will extract the area limits from the disk-header-label and supply these to the main program, thus making the area-control card unnecessary. This operation can be performed for up to four RDLIN cards. Additional RDLIN cards can be used for label checking only.

In addition to the area definition supplied by the header-label routine, the program deck can be followed by as many area-control cards as desired.

The RDLIN card is punched as follows if limits are to be extracted from the label.

Columns	Contents
1	The number of the disk drive being used
2	blank
3	\$ (indicates that limits are to be gotten from label)
4	M or L to indicate move or load mode
5, 6	Not used Note: These columns correspond to columns 15 and 16 of the area-control card. If drive address reference is to be punched into the area-control card for use when reloading with the card-to-disk program, punch reference number in column 5.
7	Track format blank = 20-sector 1 = track record
8-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification
40-80	Not used

If the disk-header-labels are to be only checked and printed, with area limits supplied in area-control cards, the routine operates in one of three ways, depending on the contents of column 21-66 of the RDLIN card.

1. If the RDLIN card contains the file-sequence number in columns 51-54, the program compares the following fields in the RDLIN card to the corresponding fields in the header labels on the track:

Columns	Contents
21-24	File Retention Period
25-29	Creation Date
30-39	File Identification
40-44	File Serial Number
45-49	Pack Serial Number
51-54	File Sequence Number

If columns 55-66 are punched, these fields are compared to the limits in the label.

2. If the RDLIN card contains no file-sequence number in columns 51-54, but does contain the file identification in 30-39, the file identification is compared to the corresponding field of the header labels. Again, if the limits are punched, they too are compared.

3. If the RDLIN card contains no file-sequence number and no file-identification field, the limits must be punched. The header labels of all files within the limits are printed.

In addition to one of the foregoing combinations, the RDLIN card must contain the drive number in column 1 and the entry RDLIN in columns 16-20.

The date card is not required with this program.

OUTPUT FORMAT

The number of cards required for one track depends on the mode and format used to write the track. In the move mode, a 20-sector track contains 2,000 character positions for records, plus twenty 6-digit addresses. These are punched into 31 cards. A 1311 track record contains 2,980 characters and one 6-digit address, requiring 43 cards. A 1301 track record contains 2543 characters and one 6-digit address, requiring 37 cards.

The number of cards required in the load mode depends on the number of word marks encountered.

Move Mode

<u>Columns</u>	<u>Contents</u>
1-5	These columns contain the five high-order digits of the address of the track being unloaded. The low-order zero is dropped. This is always the address of the zero sector of the track, regardless of the addresses actually written and regardless of the track format. This address is repeated in each card required for a track.

6-8	Card sequence number. This count is reset to 001 for each track.
9,10	Number of characters from disk storage punched in this card. In the move mode, this is always 70, except for the last card from a track.
11-80	Data from disk storage. This includes both addresses and records.

Load Mode

<u>Columns</u>	<u>Contents</u>
1-10	Punched as in the move mode.
11-30	Ten 2-column fields that indicate the card columns with which word marks are associated.
31-80	Data from disk storage; either 50 characters or 10 fields with word marks.

Example

In this example, a file of records is loaded sequentially into 1311 or 1301 disk storage in 20-sector format. The records to be punched are on the first track of the first unit in the system. Part one of Figure 5 shows the first three records as they appear in disk storage in the move mode. Thirty-one cards are required for the one track. The first four cards are shown in part two of Figure 1.

In Figure 6, the records are reduced to 90 characters and written in the load mode. The number of cards required in the load mode depends on the number of word marks in the records.

Sector	Item No.	Description	Sell	Cost	Unit Meas.	Std. Re-ord.	Ld. Tm.	Demand	Bal. O.H.	On Ord.	Min. Bal.		
1	67		2021	2526	3031	3334	3839	4142	7778	8283	8788	9293	100
000000	011234	BELT SANDER	44500	34500	EA	00001	120	0001 0002 0002 0003 0002 0001 0004 0006 0007	00020	00020	00060		
000001	021248	BOILED LIN OIL	03450	02455	GAL	00004	020	0020 0024 0035 0028 0045 0039 0052 0042 0070	00050	00000	00045		
000002	031256	BRUSH CLEANER	01250	00755	LB	00016	015	0070 0065 0092 0073 0082 0095 0092 0096 0101	00250	00000	00200		

Diagram illustrating the mapping of data from a disk to a card. The diagram shows three overlapping horizontal bars representing data segments. The top bar is labeled "3rd Record" and contains the hexadecimal string: 0000000470bb000002031256BRUSHbCLEANERb0125000755Lb0001601500700065009200730082. The middle bar is labeled "2nd Record (cont.)" and contains: 000000037055GAL00004020002000240035002800450039005200420070000500000000045bbbbbbb. The bottom bar is labeled "1st Record (cont.)" and contains: 00000002701000400060007000200002000060bbbbbbbb000001021248BOILEDbLINbOIL03450024. Below these, a larger bar represents the "70 Characters from Disk" and contains: 00000000170000000011234BELTbSANDERbbb4450034500EA0000112000010002000200030002000. On the left, vertical dashed lines indicate "Track Address" (000000), "Sector Address" (00170), and "Card Sequence Number" (011). The bottom of the diagram is labeled "1 56 891011" and "80".

21

Disk Storage Layout

Sector	Item No.	Description	Sell	Cost	Unit Meas.	Std. Re-ord.	Ld. Tm.	Demand												Bal. O.H.	On Ord.	Min. Bal.	
1	67	2021	2526	3031	3334	3839	4142	7374	7879	8384	8889	90											
000000	011234	BELT SANDER	44500	34500	EA	00001	120	0002	0002	0003	0002	0001	0004	0006	0007	00020	00020	00060					
000001	021248	BOILED LIN OIL	03450	02455	GAL	00004	020	0024	0035	0028	0045	0039	0052	0042	0070	00050	00000	00045					
000002	031256	BRUSH CLEANER	01250	00755	LB	00016	015	0065	0092	0073	0082	0095	0092	0096	0101	00250	00000	00200					

Output Format in Load Mode

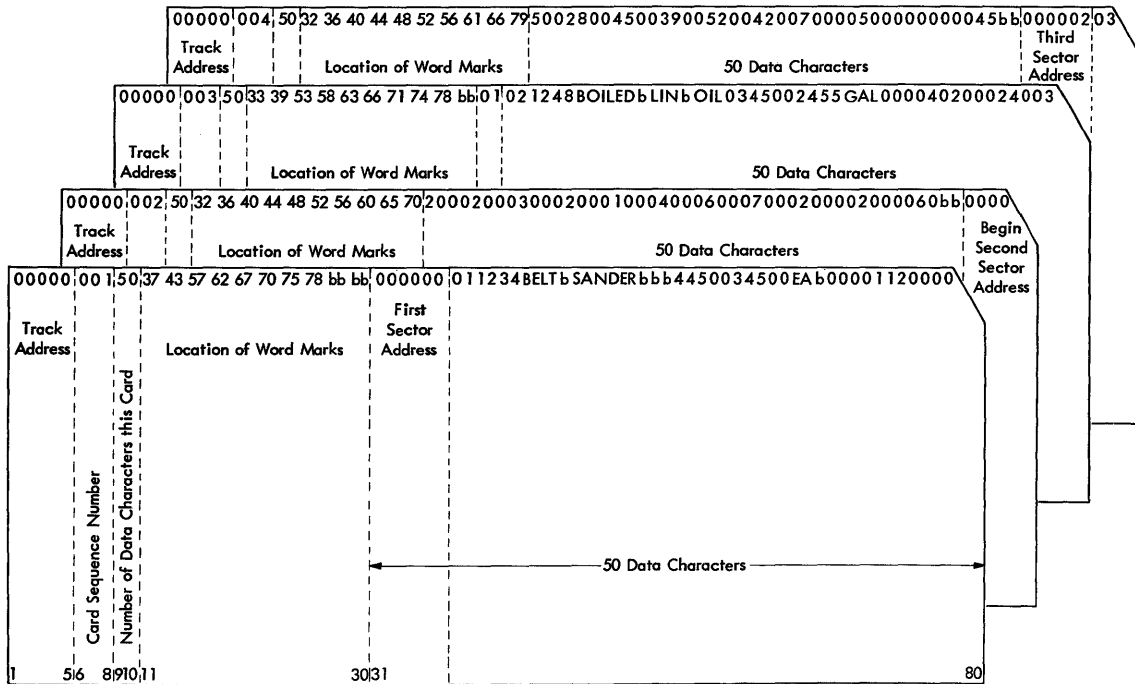


Figure 6. Disk-to-Card Example -- Load Mode

The card-to-disk program enables the user to reload into disk storage the data that was punched into cards by the disk-to-card program described previously. The data can be loaded into a disk pack or module that has any previous addressing structure, referenced to any disk drive or module. The addresses written are the same as those punched. The data is reloaded into the same relative location as that from which it was punched.

The program can reload all the data unloaded by the disk-to-card program, or it can reload selected portions of that data. The smallest area that can be reloaded is a single full track. If the user wishes to reload only selected areas, he must remove the unwanted card records from the deck before loading the program. New control cards must be supplied reflecting the change.

The card-to-disk program can be used to reload card records produced by the disk-to-card program, or card records with that same format (see Figures 5 and 6). If the program is used to perform the initial loading of a file, it is recommended that the cards be punched with each sector beginning in a new card. Thus, in the move mode, columns 9 and 10 of the first card of the two required for each sector would contain 70. Columns 9 and 10 of the second card would contain 36.

AREA-CONTROL CARDS

Each 20-column field used to define an area to be reloaded is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Mode - M or L
2-7	Lower limit
8-13	Upper limit
14	Disk drive or module being used
15	Disk drive or module to which addresses are referenced. <u>Note:</u> This number can be prepunched in the control cards used with the disk-to-card program. That program does not check this column.
16	Not used
17	Track format blank = 20-sector 1 = track record
18-20	Not used

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the card-to-disk program is used to insure that all files within the

disk area being loaded are expired and, if so, to enter a header label for the card file.

The user punches in one or more RDLIN cards the limits of the area(s) to be loaded. The program searches the label track for labels that pertain to any portion of the limits provided. Labels are checked for files within the same relative limits with addresses referenced to other disk drives. For example, if the limits given were 040040-055019, the labels on the pack would be searched for files within the following limits:

040040-055019
060040-075019
080040-095019
000040-015019
020040-035019

If the retention cycle of any label found by this search has not expired, the program prints a message to that effect along with the unexpired label(s) and then halts.

If all labels pertaining to the area are expired, they are deleted by blanking the file-identification field. All labels so deleted are printed out. The routine then enters the new label, as punched in the RDLIN card, on the track.

RDLIN CARDS

A RDLIN card must be supplied for each area to be loaded. Each card must contain the limits of the area to be checked for active header labels. The header label written will contain these limits and any other fields punched in the RDLIN card with the exception of the pack serial number. The card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	The number of the disk drive being used
2-15	Not used
16-20	RDLIN
21-29	Optional. If punched, these fields are written in the header label
30-39	File identification
40-50	Optional. If punched, these fields are written in the header label.
51-54	Not used (old pack serial number is retained)
55-60	The lower limit of the area
61-66	The upper limit of the area
67-80	Not used

A date card must be supplied to this routine.

COPY-DISK PROGRAM

The copy-disk program is used to copy a file of disk records from one area of disk storage to another. Transfer of information can be between:

- 1311 to 1301
- 1311 to 1311
- 1301 to 1311
- 1301 to 1301

The program can be used to perform the following operations.

1. Copy all or part of the contents of one unit onto another.
2. Relocate a file from one area of disk storage to another area on the same or another disk unit.
3. Relocate into the same area on the same pack or module while changing the address references.

The purpose of the first operation is similar to that of the disk-to-card and card-to-disk programs. Normally, when large files of master records must be duplicated, the copy-disk program is better suited to the job than are the card programs. If small areas are being unloaded, the card programs have certain advantages. However, large areas of disk storage or entire disk packs or modules can be copied much more quickly with the copy-disk program and can be stored more easily.

The second operation is performed when it is desired to relocate data in disk storage in order to:

1. reduce seek time for a file by placing it nearer the home arm position, or
2. make room for other files.

It is sometimes necessary to change the addressing reference on a disk pack or module. The third operation might be performed to ready a multi-pack input file for a program that requires consecutive disk addresses throughout the file.

When copying information written in the track-record mode, from a 1311 to a 1301, the last 437 characters (421 in load mode) of the longer 1311 track-record are lost. These positions are padded with blanks when copying from 1301 to 1311.

AREA-CONTROL CARDS

If a file is being copied into the same relative position in another disk pack, the area-control information can be supplied to the program by the header-label routine. If area-control cards are used, they are punched in two different ways, depending on whether data is being relocated or copied into the same relative location on another pack while keeping the addresses of the pack or module being copied.

Copy Without Relocation

If data is to be copied from one disk pack or module onto another while retaining the same relative location within the pack or module and keeping the addresses of the pack or module being copied, each area-control card can define as many as four areas. Each area is defined in a 20-column field punched as follows:

<u>Columns</u>	<u>Contents</u>
1	M or L to indicate move or load mode
2-7	Lower limit address of area to be copied
8-13	Upper limit address of area to be copied
14	The number of the disk drive or module on which the data being read is located
15	The number of the disk drive or module on which the data being written is located
16	The number of the disk drive or module to which the addresses of the pack or module being written are referenced
17	This column contains a 1-punch if data is written in track record format
18-20	Not used

Copy With Relocation

If a disk file is to be relocated in disk storage, an area-control card must be supplied. The area information can not be supplied by the disk-header-label routine. Only one area is defined in each area-control card used for this operation. Each card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1-17	These columns are punched exactly as when copying without relocation
18-20	REL punched in these columns indicates that the file is to be relocated
21-26	Lower limit address of output area
27-32	Upper limit address of output area
33-80	Not used

DISK-HEADER-LABEL ROUTINE

The header-label checking procedure followed by the copy-disk program depends on whether a file is being relocated or copied into the same relative disk locations.

If the file is not being relocated, the routine searches the input pack for the header label matching the File Identification field of the RDLIN card. If this label is found, it is printed out. The routine then extracts the area limits from the input-header label and uses these to search for unexpired files on the output pack. Labels are checked for files within the same relative limits with addresses referenced

to other disk drives. If all files within those limits are expired, the routine deletes and prints the expired labels. The input-header label is written on the output pack with only the pack serial number changed.

If the file is to be relocated, two RDLIN cards are used. The first card is used to read the input-header label. The routine searches for and prints the header label requested.

The second card is used for the output-header label. The routine searches the output pack for un-expired files within the limits given in the card. Expired labels are deleted, and the information in the RDLIN card is written as a new label. Again, the old pack serial number is retained, regardless of whether this field is punched in the RDLIN card.

RDLIN CARDS

If the file is not to be relocated, a single RDLIN card is used. The card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	The number of the input disk drive.
2	The number of the output disk drive.
Columns 3-7 are used only if the header label routine is to supply area control information to the main program.	
3	\$
4	Mode - M or L
5	Blank
6	The number of the disk drive to which the addresses on the output pack are referenced.
7	Track format blank = 20-sector 1 = track record
8-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification of file to be copied
40-80	Not used

If the file is to be relocated, two RDLIN cards are used. The cards are punched as follows.

Input-Header-Label Card

<u>Columns</u>	<u>Contents</u>
1	The number of the input disk drive
2	The letter N
3-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification of file to be copied
40-80	Not used

Output-Header-Label Card

<u>Columns</u>	<u>Contents</u>
1	The letter N
2	The number of the output disk drive
3-15	Not used
16-20	RDLIN
21-29	Optional. If punched, these fields are written in the header label.
30-39	File identification to be written in label
40-54	Optional. If punched, these fields are written in the header label
55-60	Lower limit of relocated file
61-66	Upper limit of relocated file
67-80	Not used

A date card must be supplied to this routine.

PRINT-DISK PROGRAM

Most disk-storage operations that produce printed output require specific programs to select the portions of each record that are to be printed, arrange these fields in a particular format, accumulate totals, etc. However, it is frequently desirable to obtain a printed copy of an area of disk storage exactly as it is written on the disk. This is especially true during program testing. The print-disk program produces such a copy for any number of areas without programming effort on the part of the user.

The print-disk program can print tracks written in either the move or load mode and in either 20-sector or track-record format. If desired, special characters that are not printable with the standard 48- and 52-character set can be replaced with a printable alphameric or special character.

PRINT FORMAT

An index, which aids in determining character positions, is printed at the top and bottom of each page. In the 20-sector format and the track-record format, an M or L (to indicate the mode used to write the disk) and the address of the record, are printed before the printing of the record.

When printing areas written in the load mode, the user has the option of printing an indication of word marks. A word mark is indicated by printing a one (1) under the character with which it is associated.

AREA-CONTROL CARDS

If a labeled file is to be printed, the area-control card is not necessary (see Disk-Header-Label Routine). If area-control cards are used, each 20-column field used to define an area to be printed is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Mode - M or L
2-7	Lower limit
8-13	Upper limit
14	Disk drive or module being used
15	Character to be substituted for unprintable characters
16	Word mark option for load mode: blank = print word marks 1 = suppress word mark printing
17	Track format blank = 20-sector 1 = track record
18-20	Not used

DISK-HEADER-LABEL ROUTINE

The disk-header-label routine for the print disk program is used to insure that the proper file is on-line and, if so, to print its header label. If desired, the routine will extract the area limits from the disk label and supply these to the main program, thus making the area-control card unnecessary. This operation can be performed for up to four RDLIN cards. Additional RDLIN cards can be used for label checking only.

In addition to the area definitions supplied by the header label routine, the program deck can be followed by as many area-control cards as desired.

The RDLIN card is punched as follows if limits are to be extracted from the label.

<u>Columns</u>	<u>Contents</u>
1	The number of the disk drive being used
2	Blank
3	\$ (indicates that limits are to be gotten from label)
4	M or L to indicate move or load mode
5	Character to be substituted for unprintable characters
6	Word mark option for load mode: blank = print word marks 1 = suppress word mark printing
7	Track format blank = 20-sector 1 = track record
8-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification of file to be printed
40-80	Not used

If the header labels are to be only checked and printed, with area limits supplied in area-control cards, the routine operates in one of three ways, depending on the contents of columns 21-66 of the RDLIN card.

1. If the RDLIN card contains the file-sequence number in columns 51-54, the program compares the following fields in the RDLIN card to the corresponding fields in the header labels on the track:

<u>Columns</u>	<u>Contents</u>
21-24	File Retention Period
25-29	Creation Date
30-39	File Identification
40-44	File Serial Number
45-49	Pack Serial Number
51-54	File Sequence Number

If columns 55-66 are punched, these fields are compared to the limits in the label.

2. If the RDLIN card contains no file-sequence number in columns 51-54, but does contain the file identification in 30-39, the file identification is compared to the corresponding field of the header labels. Again, if the limits are punched, they too are compared.
3. If the RDLIN card contains no file sequence number and no file-identification field, the limits

must be punched. The header labels of all files within the limits are printed.

In addition to one of the above combinations, the RDLIN card must contain the drive number in column 1 and the entry RDLIN in columns 16-20. Columns 2-15 must be blank.

The date card is not required with this program.

DISK-RECORD-LOAD PROGRAM

The disk-record-load program is designed for small amounts of data. The program loads addresses, single records, parts of records, or entire tracks into 1311 or 1301 disk storage. The program reads the specified track-record, sector, or full track with addresses into core storage, and inserts the data from cards into the desired locations. This program can be used in program testing to set up a sample file of test records. Or, it can be used to correct or change existing files. The program can change either addresses or data on a track.

If one record is placed in several locations on a track, the program can change all the sectors with the same address as that specified.

This program does not require separate area-control cards. Instead, each input card contains both the data to be inserted and the control information. The first 20 columns of each card contain the required control information. The data to be inserted is punched in the remaining 60 columns.

In the move mode, data is punched beginning in column 21 and continuing for as many columns as necessary. If more than 60 columns are required, additional cards must be punched, each containing the control information pertinent to the portion of the field or record contained in that card.

If the disk record is written in the load mode, only one field is punched in each card. Again, this field is punched beginning in column 21. The user must specify in the control portion whether the first character of the field is to contain a word mark.

If the change is to be made to a disk address, the new six-digit address is punched in the data portion of the card. A dollar sign in column 11 indicates that the change is to be made to the address specified. Columns 8-10 and 12-13 can be blank in this case.

INPUT FORMAT

Each card is laid out as follows:

Column	1	Contains an M or an L to indicate either move or load mode
Columns	2-7	This six-column field contains the address of the track or sector to be operated on
Columns	8-11	This 4-column field specifies the location of the data within the record (high-order position within either the sector or the track record, depending on which format is used (first position is 0001). If changing address, punch \$ in column 11, and columns 8-10 must be left blank.
Columns	12-13	These columns specify the length of the data field (maximum 60). These columns must be left blank if changing addresses.
<u>Note:</u> High-order zeros must be punched in field 8-11 and 12-13		
Column	14	Specifies the disk drive or module being referenced
Columns	15, 16, and 17	are used to specify the user's choice of three options available with this program:
Column	15	This column is used to specify that the program shall: a. print an analysis of this card and halt before processing (blank) b. print analysis and continue processing without halting (1-punch) c. bypass print and halt of analysis routine except in event of error (2-punch)
Column	16	A 1-punch indicates that the program is to scan the entire track and alter all records with identical sector addresses
Column	17	If in the load mode, a 1-punch in this column indicates that a word mark is to accompany the first character of data.
Column	18	A 1-punch in this column indicates the track-record format.
Columns	19-20	Must be left blank
Columns	21-80	Contain the data or address to be inserted.

DISK-HEADER-LABEL ROUTINE

The header-label routine used with the disk-record-load program does not generate area-control information but otherwise operates in exactly the same way as the routine used with the print-disk program.

The disk-label program performs all necessary maintenance operations on the label track of a disk pack. The program is used to:

1. Set up the initial header-label track on a disk pack.
2. Remove the entire label track by clearing it and restoring the addresses to the original range.
3. Enter a new label.
4. Delete an existing label.
5. Make changes to an existing label.
6. Print labels.
7. Punch and print labels.

RDLIN CARDS

RDLIN cards are used as input to this program. Normally, the standard RDLIN card shown in Figure 2 is used. When it is necessary to operate on fields not contained in the standard RDLIN card, an expanded format is used (Figure 7). The expanded format consists of two cards. The first five columns of the first card are punched as in the standard RDLIN card. Columns 6-9 of the first card contain XPN1. Columns 1-5 of the second card are left blank, and columns 6-9 of the second card contain XPN2.

Columns 16-20 of both cards contain RDLIN.

The header label fields are punched, as shown in Figure 7, in the same order as they appear in the 100-character disk label.

SPECIFIC OPERATIONS

Set Up Initial Label Track (NEW)

When setting up the initial header-label track on a previously unlabeled disk pack, the program performs the following operations:

1. Change the addresses on the label track from 0X9980-0X9999 to 000180-000199.
Note: The address of the zero-sector (0X9980) must be present on the track at the beginning of this operation.
2. Clear all 20 sectors to blanks.
3. Enter the label identifier (1HDRb) in the first 19 sectors and the pack serial number in all 20 sectors. The standard, single-card RDLIN card

must be used for this operation. It is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-4	NEW
5	Disk drive to which the addresses on the pack are referenced
6-15	Not used
16-20	RDLIN
21-44	Not used
45-49	Pack serial number
50-80	Not used

The date card is not used with this operation.

Restore Normal Addresses (RNA)

This operation clears the entire label track to blanks and restores the addresses to the range: 0X9980-0X9999.

The standard, single-card RDLIN card must be used for this operation. It is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-4	RNA
5	Disk drive to which the addresses on the last track are to be referenced
6-9	Blank
10-15	Not used
16-20	RDLIN
21-80	Not used

The date card is not used with this operation.

Enter a New Label (ENTR)

This operation is used to enter a complete new header label. The program first searches the label track to insure that there are no unexpired header labels that pertain to files within the limits given in the RDLIN card. Labels are checked for files within the same relative limits with addresses referenced to other disk drives. Any expired labels pertaining to this area are printed and then deleted. The program then enters the new label from the RDLIN card. The old pack serial number is retained, regardless of whether it is punched in the RDLIN card.

Either the standard or the expanded RDLIN card format can be used. If the standard format is used, all fields not contained in the RDLIN card are blanked.

The RDLIN card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-5	ENTR
6-9	Blank, XPN1, or XPN2
10-15	Not used
16-20	RDLIN

The rest of the fields in the RDLIN card, with the exception of the pack serial number (columns 45-49), are entered in the label. If any fields are left blank, blanks are written in the corresponding positions of the disk label. The lower and upper limits must be punched.

The date card is required for this operation.

Delete Labels (DELT)

This operation is used to delete either all of the expired header labels pertaining to files within specified limits or to delete any specific label. If an unexpired label is to be deleted, it must be deleted by specific file identification. All labels deleted are printed.

The standard, single-card RDLIN card is used for this operation. The RDLIN card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-5	DELT
6-9	Blank
10-15	Not used
16-20	RDLIN
21-29	Not used
30-39	File identification, if a specific header label is to be deleted. Otherwise, these columns are blank
40-54	Not used
55-66	Lower and upper limits, if all labels within limits are to be deleted. Otherwise, these columns are blank. If limits are to be used, columns 30-39 must be blank.
67-80	Not used

The date card must be supplied for this operation if all labels within specified limits are to be deleted.

Change a Label (CHNG)

This operation is used to make changes to individual fields within a header label. The program finds and prints the header label specified in the file-identification field of the RDLIN card. All non-blank fields in the RDLIN card are then substituted for the corresponding fields in the header label. The changed label is then printed and written back onto the label track.

Either the standard or the expanded RDLIN card format can be used. The card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-5	CHNG
6-9	Blank, XPN1, or XPN2
10-15	Not used
16-20	RDLIN
30-39	File identification (must be punched)

All other fields are entered in the corresponding fields of the label if punched.

The date card is not used with this operation.

Print Labels (PRNT)

This operation prints either all of the header labels on a disk pack or only specifically requested labels. The standard, single-card RDLIN card is used with this operation. The RDLIN card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-5	PRNT
6-9	Blank
10-15	Not used
16-20	RDLIN
21-23	ALL, if all header labels are to be printed
24-29	Not used
30-39	File identification, if specific label is to be printed
40-80	Not used

The date card is not used with this operation.

Punch and Print Labels (PNCH)

This operation punches and prints either all of the header labels on a disk pack or only specifically requested labels. The labels are punched in either standard or expanded RDLIN card format. The standard, or single-card is used with this operation. The RDLIN card is punched as follows:

<u>Columns</u>	<u>Contents</u>
1	Disk drive being used
2-5	PNCH, if labels are to be punched in the standard RDLIN card format PNCX, if labels are to be punched in the expanded RDLIN card format
6-9	Blank
10-15	Not used
16-20	RDLIN
21-23	ALL, if all header labels are to be punched and printed
24-29	Not used
30-39	File identification, if specific label is to be punched and printed
40-80	Not used

The date card is not used by this operation.

[illegible]

Figure 7. Expanded RDLIN Card Format

TIMINGS

The processing times shown in Figure 8 are approximate. They are given to facilitate estimating total job times. The times shown are those required to process one full pack or module with these configurations.

1401-1311	1403 Printer Model 3
1440-1301	1442 Reader Model 1, 1443 Printer Model 2
1440-1311	1442 Reader Model 2, 1443 Printer Model 1

1460-1301 1403 Printer Model 3

1460-1311 1403 Printer Model 3

The time required to copy areas of disk storage from one unit to another depends on the relative positions of the index points on the two disk units. The operation is bound by its input and output, regardless of the speed of the processing unit. The time required to copy 20,000 sectors of disk will vary between 3.5 and 7 minutes, regardless of the type of disk unit.

Disk-to-Tape	729-2 High Density	729-2 Low Density	7330 High Density	7330 Low Density	7335
1401-1311	2.6 min	4.5 min	5.3 min	8.6 min	
1440-1301					38.4 min
1440-1311					3.9 "
1460-1301	17.3 "	35.3 "			
1460-1311	2.6 "	3.9 "			
Tape-to-Disk					
1401-1311	4.0 "	5.9 "	5.3 "	8.6 "	
1440-1301					40.0 "
1440-1311					5.3 "
1460-1301	29.7 "	46.5 "			
1460-1311	3.7 "	5.2 "			

Disk-to-Card (Note 1)	Move Mode (in minutes)	Load Mode (in minutes)
1401-1311	128 min	235 min
1440-1301	6050 "	7617 "
1440-1311	353 "	441 "
1460-1301	1283 "	1976 "
1460-1311	128 "	197 "
Card-to-Disk (Note 1)		
1401-1311	44 "	54 "
1440-1301	1000 "	1250 "
1440-1311	97 "	117 "
1460-1301	443 "	533 "
1460-1311	44 "	53 "
Print Disk (Note 2)		
1401-1311	40 "	75 "
1440-1301	250 "	416.6"
1440-1311	145 "	292 "
1460-1301	250 "	416.6"
1460-1311	25 "	41.6"

Note 1: The time required to load or unload a disk pack in the load mode using the card programs depends on the frequency of word marks in the file. The time given here is for the best case, where each card contains 50 data characters.

Note 2: The special-character substitution option requires an additional 25 minutes for systems with a 1311 attached or an additional 250 minutes for a 1301, if the print storage feature is not available.

Clear Disk	Same Address	New Address
1401-1311	5.2 min	3.1 min
1440-1301	58.0 "	33.0 "
1440-1311	5.8 "	3.2 "
1460-1301	56.0 "	34.0 "
1460-1311	4.4 "	3.1 "

Figure 8. Timings

Disk Utility Programs Specifications
IBM 1401, 1440, 1460 (with IBM 1301 and 1311)
C24-1484-3

- | | <i>Yes</i> | <i>Satisfactory</i> | <i>No</i> |
|--------------------|--------------------------|--------------------------|--------------------------|
| • Is the material: | | | |
| Easy to read? | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Well organized? | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Fully covered? | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Clearly explained? | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Well illustrated? | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

- How did you use this publication?

As an introduction to the subject	<input type="checkbox"/>
For additional knowledge of the subject	<input type="checkbox"/>

- Which of the following terms best describes your job?

<i>Customer Personnel</i>		<i>IBM Personnel</i>	
Manager	<input type="checkbox"/>	Customer Engineer	<input type="checkbox"/>
Systems Analyst	<input type="checkbox"/>	Instructor	<input type="checkbox"/>
Operator	<input type="checkbox"/>	Sales Representative	<input type="checkbox"/>
Programmer	<input type="checkbox"/>	Systems Engineer	<input type="checkbox"/>
Trainee	<input type="checkbox"/>	Trainee	<input type="checkbox"/>
Other _____		Other _____	

- Check specific comment (if any) and explain in the space below:

(Give page number)

- | | |
|---|---|
| <input type="checkbox"/> Suggested Change (Page) | <input type="checkbox"/> Suggested Addition (Page) |
| <input type="checkbox"/> Error (Page) | <input type="checkbox"/> Suggested Deletion (Page) |

Explanation:

fold

fold

FIRST CLASS
PERMIT NO. 387
ROCHESTER, MINN.

BUSINESS REPLY MAIL
NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

POSTAGE WILL BE PAID BY . . .

IBM Corporation
Systems Development Division
Development Laboratory
Rochester, Minnesota 55901

Attention: Programming Publications, Dept. 425

fold

fold

IBM 1401/1440/1460 Printed in U. S. A. C24-1484-3

IBM

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



International Business Machines Corporation

Data Processing Division

112 East Post Road, White Plains, N.Y. 10601