

COMPUTATION CENTER
Massachusetts Institute of Technology
Cambridge 39, Massachusetts

May 9, 1963

TO: Computation Center Staff
FROM: Richard H. Orenstein, Robert C. Daley
SUBJECT: LD EDT and DP EDT, the CTSS Disk Editors

A. INTRODUCTION

This memo outlines the mechanics and usage of the two disk editor programs, whose preliminary specifications are mentioned in CC-196. They are reviewed here.

The master file directory (MFD) starts on track 0, module^{*} 1 and is chained as is any other track, the first word being a pointer to the next track, if any, or a word count for this last track. The MFD is made up of two-word entries. The first word contains the user's problem number in the address and the user's programmer number in the decrement. The second word contains a pointer to the first track of the user file directory in the address and tag, and the user track quota in the decrement.

Starting on track 1, module 1 is the track usage table (TUT). This is a table of all tracks used or unused for all modules. The status of each track is denoted by a bit in a word. For a used track this bit is a "1", for an unused track this bit is a "0". This table is the only one not chained, but written on sequential tracks.

Each user has a user file directory (UFD) which contains a 4-word entry for each file. The first two words contain the primary and secondary name of the file. The third word contains the mode of the file in the prefix, a documentation number (a number assigned internally to each file for tracing purposes) in the decrement, and the address of the first track of the file in the address and tag. The fourth word contains the number of tracks and by the file in the address and the date the file was last used in the decrement.

The edit programs are run within the regular FMS system under the FMS monitor. The control cards and other input are data for the programs.

*Logical module numbers are used throughout this memo.

Facilities are as follows:

1. Dump tapes can be prepared singly or in duplicate, under sense switch or control card setting.
2. During reloading files are reorganized so that consecutive tracks of a file are assigned sequential track addresses.
3. The user's track quota may be changed by changing the MFD entry card.

By use of control cards:

4. Files are input to the disk or edited to the disk from binary or decimal cards.
5. Output tapes are prepared from files for off-line printing or punching.
6. Files may be dumped for off-line punching according to the date last used.
7. Cards of files dumped according to the date last used may be reloaded into the original file.
8. Files may be deleted.
9. The mode of a file may be changed.
10. A file may be copied from one user to another.
11. A summary of track usage for any or all users may be output for off-line printing.
12. Tracks may be deleted from usage in case they are consistently bad for reading and writing.
13. Dump and load time may be charged to various problem and programmer numbers.
14. Home addresses may be written on a specified number of tracks to minimize time needed during a write operation.

Both editor programs do their own time accounting via FMS subprograms that use the interval timer. Within both programs an end-of-file mark may be either the standard end of file card (7-8 punches in column 1 and "END OF FILE" in columns 16-26) or a decimal card blank in columns 1-7 with "EOF" in columns 8-12.

B. The Dump Editor

The following sense switch settings affect the program

Sense Switch 2: up-write duplicate dump tapes; down-write single dump tapes.

Sense Switch 5: up-off line printing only off-line; down off-line printing appears on line also.

Sense Switch 6: up-continuous execution of control cards; down-stop after execution of each control card (for debug purposes)

Initialization consists of setting the interval timer and the disk trap and interrupt locations and examining sense switch 2. Other sense switches are examined when the branch is encountered so their setting may be changed at any time. The first track of the master file directory is then read into core.

Control immediately passes to the section that reads and processes the control cards. Since changes in the file directories may take place, a certain order should be maintained among the control cards. The format for control cards allows a variable number of fields, each of a variable length not exceeding 6 characters. Fields are separated by commas, and the card is terminated with a blank.

The general format of the control card, to be used unless otherwise specified, is as follows:

1. Control word
2. Problem number
3. Programmer number
4. Primary file name
5. Secondary (class) file name

Below is a list of control words, their requirements and their functions. For those marked with an asterisk the time they use is charged as editor system time.

1. **STAT*** There is an option of parameters, either the word "ALL", or some problem and programmer number. If a problem and programmer number is given, then statistics are given for that user; if "ALL" is used, then statistics are given for all users. Statistics consist of the problem and programmer number, the user's track quota, the total number of tracks used, and the track address of the first track of the UFD. Following is a list of the user's files with the name, document number, the mode of the file, the address of the first track of the file, the number of tracks used by the file, and the last date that the file was used.

2. DATED. The three parameters here are integers representing the month, day and year. During the dump operation, any files (of all users) whose date last used is earlier than the date specified on the DATED card will automatically be deleted along with the production of output cards and an appropriate comment. The cards produced here are a special form of column binary with the prefix of the first word on every card being a "7". Also, instead of a loading address, these cards have a sequence number whose order must be preserved. The last card of the deck has a tag of 1 in the control word so that completeness of the deck may be verified. These cards may be reloaded onto the disk following the INPUT conventions (Part D). This card must precede the DUMP card.

3. DELETE This control card has the standard format with an optional sixth field. The file specified is deleted from the user's file directory. If the word "CARDS" appears as a sixth parameter, the prefix 7 cards, described in Section 2, will be produced off-line. This card must precede the DUMP card.

4. DUMPe. This control card causes the disk to be dumped, user by user, file by file, onto pre-assigned tapes. There is an optional parameter which may be any non-zero, non-blank character. If this character is present, single dump tapes will be written. Either this parameter or sense switch 2 down will produce single dump tapes; if either condition is sensed, dump tapes are written in duplicate.

Drives A5, A6, B5 and B6 are available for dumping. If single dump tapes are to be written, then only those drives on Channel A will be used. Dumping commences onto A5 and B5. When these reels are full they are re-wound and unloaded and A6 and B6 are used. Should these reels become full with more to dump, A5 and A5 are used again. This cycle repeats as long as necessary to complete dumping.

First on every dump tape is a 4 word file called the heading. This is merely a label used to check that tapes are mounted in the proper order at load time. The four words are

1. The type of dump (there is another form discussed later)
2. The date
3. The time
4. A sequential number given each reel.

--and are written in BCD so they may be printed for identification purposes.

The type of dump discussed here is the normal form. An emergency form (FDUMP) is available as an absolute on-line job. Its purpose and use will be described later.

The actual dump procedure begins by scanning the MFD for the first, or next entry. The first word of the entry (problem and programmer number) is saved and a new second word is formed containing the last document number in the decrement. These two words plus the four-word entry from the UFD make up a six word record that precedes each file as it is dumped on tape. For each MFD entry, the UFD is scanned and each file is dumped in turn. The files are dumped in records corresponding to each track as it is on the disk, including the chain word. Each file, preceded by its 6-word control record, is dumped until the end of the UFD. At this point end of file is written on the dump tape and the cycle repeats for the next MFD entry. If an end of tape is encountered, two ends of file are written and the tape is rewound and unloaded. To dump an entire disk module with 556 bit/inch tape density requires 2-1/2 tape reels.

If a track of the disk is unreadable, a code of 5 is placed in the prefix of the control word for the record as written on tape. The error is logged on-line. This code will be detected during loading but will not stop normal operating procedures.

The document number is a cyclic integer (1-32767) uniquely assigned to each file as it is added to the file directory. This assures a unique number for each file even though there may be name duplication. It was thought that such an identification would be needed for each file in case of retrieval problems (a record for a file is unreadable from tape or disk). However, no use of the document number has been incorporated into the present version of the dump-or load-editor.

5. PRINT This control card has the general format. It causes the specified file to be written on the FMS output tape for off-line printing. The file is assumed to be 14 words per line unless otherwise designated. A blank word is inserted at the beginning of each line to insure single spacing. If the first word of a file is a line mark, this will then specify the number of words per line. The line mark is a word whose prefix and decrement are all 1's, and whose address contains the number of following words to be printed. This then assumes that the first character of the first word after each line mark will be the printer control character for spacing. Only files written on the disk through special subroutines will have this line-mark feature.

6. DPUNCH This is the same as print except that the file is written on the FMS punch-output tape. Line-marked files are also recognized.

7. BPUNCH This control card has the general format. It causes the specified file to be written on the FMS punch-output tape for off-line punching. The file is written on tape for punching 28 words per card. No provision is made for 7-9 punches or check sums which are assumed to be included in the card image.

8. 7PUNCH This control card, with general format, causes prefix-7 cards to be punched for the specified file without deletion of the file.

9. CHMODE This control card, with the general format, changes the mode of the specified file. The new mode is placed in an additional field on the control card. If the new mode is temporary, the user track count is adjusted accordingly. The mode is specified by a single digit, as follows:

- 0 = temporary
- 1 = permanent
- 2 = read-only class 1
- 3 = read-only class 2

10. COPY This control card causes a specified file of one user to be copied for another user so that each user will have an individual copy. The general format is used to specify where the file is going and what it is to be called. Four additional fields specify where to get the file. They are the problem and programmer number, and the two file names there. The mode of the file is retained during this operation.

11. CHARGE* This control card will cause system edit time to be charged to the problem number and programmer number specified by the two fields on this card. If this card does not appear within the deck, time will be charged as editor system time.

12. CLOSE* This control card signals the end of the data. It has no parameters. Upon reading this card, a final time card is produced for the total system edit time. Control is then returned to the monitor system.

The Fast Dump

Only sense switch 2 affects this program; it is used to control the number of dump tapes written. If it is down then single dump tapes will be written. The usage of dump tapes here is the same as with the DUMP facility in the dump editor. That is, drives A5 and B5 will be used first. When they are full drives A6 and B6 will be used. Should reels on these drives become full, A5 and B5 will be used again. If single dump tapes are to be written only the drives on channel A will be used.

This program immediately writes the heading file on the first tapes used. The heading is exactly that described in section B.4 except that the first word is "FDUMP." The remaining three words are then the date, time and sequence number.

The track usage table is then read into core and written as one record onto tape. Then the table is scanned, and each occupied track is dumped as one record (466₁₀ words). When all the occupied tracks for the module are dumped, the process continues for the next module and so on until no more modules exist.

At this point a comment is printed to the operator and the program halts.

D. The Load Editor

The following sense switch settings effect this program

Sense Switch 1: up-write format tracks for all cylinders; down-do not write any format tracks.

Sense Switch 2: up-read from duplicate dump tapes; down-read from single dump tapes.

Sense Switch 5: up-off-line printing only off-line; down-off-line printing appears on line also up-continuous execution of control cards down-halt after execution of each control card.

Sense Switch 6: up-continuous execution of control cards; down-halt after execution of each control card.

Initialization consists of setting the interval timer for time accounting, testing sense switch 2 for tape setup, and testing sense switch 1. According to the setting of sense switch 1, format tracks are either written or not. The first data cards of the load editor FMS job should be the MFD cards unless a fast load is requested. This latter case will be discussed in a following section. Each entry in the MFD appears on a single card which has three 6 column fields. All the information must be left justified within each field. The fields contain

1. The problem number (Cols. 1-6)
2. The programmer number (Cols. 7-12)
3. The user's track quota (Cols. 13-18)

Columns 19-80 are ignored.

The MFD entry cards are read and the MFD is written on the disk track by track until an end-of-file signal is sensed. After complete assembly of the directory, the first track is read into core to facilitate searching.

Now control cards are read and processed. The format for control cards is the same as for the dump editor -- a variable number of fields, each of a variable length not exceeding 6 characters, separated by commas, terminated by a blank.

The following list of control cards are available in the load editor. In this list an asterisk, as before, will indicate time to be charged as editor system time.

1. **LOAD*** This control card causes the disk to be loaded from tapes used for dumping the disk by the dump editor. There is an optional parameter which may be any non-blank, non-zero character. If there is such a parameter, only a single tape will be examined for heading. If no parameter is present, duplicate tapes will be assumed to be mounted. If sense switch 2 is down, single tapes will be assumed, and this will override no parameter on the LOAD card.

The load section first checks the heading type and sequence number. If a discrepancy is detected, the operator is informed by some message on-line, and a program stop. At this point there is an option depending on the contents of the console keys. If the keys contain zero and start is depressed, the error will be ignored. If something non-zero is put in the keys and the start key depressed, the tapes in error will be rewound and re-examined for correct headings. When the headings are verified, loading of files, etc. commences.

The end of file mark following the heading is skipped and the first or next 6-word record is examined. The end of file signal shows a new MFD entry follows. The user is looked up in the MFD that has been assembled and if he is present loading continues. From the 6 word record for the files, the module number for the file is determined, and the first track for the file is reserved. A check is made to assure that by loading the file the user's track quota will not be exceeded. The file is loaded and an entry is made in the user's file directory. Each file is loaded in the same manner until an end of file mark is sensed. Then the user's file directory is written on the disk and the process continues for the next user. An end signal - a special coded user number-signals no more entries follow and loading is terminated.

If the error code for an unreadable track is detected in a record from tape (a prefix of 5₀ in the first word), the record is loaded but a comment is put out on-line telling which file has the error. If a record is unreadable from tape, the file of which the record is a part is deleted with a comment on-line.

During the load process when duplicate tapes are in use, the tape on channel A is used primarily. For each record (excepting the heading) that is read, a record on the duplicate tape is skipped. If a redundancy occurs, then the record on the duplicate tape is read. If that tape also has a redundancy check for the record, the tapes are alternatively tried 9 more times. If still unsuccessful, the record is skipped and considered bad.

2. EDIT This facility provides a means of loading a binary deck to form a core image of a program which is then written as one record on the disk. The format for the control card is the general one with an additional field containing an integer specifying the mode of the file to be created as described under CHMODE. The control card is followed by a column binary record card for the edit deck. This record card has a first word with a prefix of 1. The third word of the card contains in the decrement the last address used by the program. The following cards must be absolute column binary and are checked for correct check sums, etc. The absolute cards are loaded until an end of file signal is reached. At this point the file is written on the disk, and then entered into the user's file directory.

3. INPUT This control card has the general format plus an extra field specifying the mode of the file to be created as described under CHMODE. The deck of cards immediately following the control card is put on the disk as a single file. Binary and decimal cards may be mixed. If the card is binary then 28 words are read if it is decimal then 14 words are read. When the end of file is read, the file is added to the user's file directory.

The cards produced by a DELETE card or a DATED card in the DUMP editor should be reloaded by this control card. The prefix of 7₈ is recognized and then the file is recreated from these cards. The sequence numbers on these cards must be in order or an error will occur.

4. DELTRK* This control card causes specified tracks to be marked off as used in the track usage table. It is thereby a means of not using certain tracks, since these tracks will be pseudo-assigned. The format for this card is as follows:

- a. After the control word, the next field contains the logical module number for the tracks to be deleted.
- b. The track address to be deleted.
- c. An optional field. If blank, only the single track mentioned in field b will be deleted. If non-blank then the number in the field specifies the number of sequential tracks to be deleted, starting at the track address specified in field b.

5. DELSUR* This control card performs the same function as DELTRK except to a surface of the disk. There are only two parameters, the first specifies the module number and the second specifies the surface to be deleted. The entire 250 tracks of the surface will be marked in the track usage table.

6. U-1* This control card causes home addresses to be written on the specified number of tracks. There are two parameters, the first specifies the module number, the second specifies the number of tracks to be written with home addresses. The count is made by including all tracks used up to this point as having home addresses, so that on exit from this section no more than the specified number of tracks have home addresses written on them.

7. CHARGE* This control card causes the system edit time to be charged to the problem and programmer number specified by the two fields on this card. If this card does not appear in the deck, time will be charged as editor system time.

E. Typical Data Deck Setup for Loading

