



Systems Reference Library

Sort 1 for IBM 1401 – Specifications

Presents specifications for Sort 1, a generalized sorting program for use on an IBM 1401 Data Processing System equipped with a minimum of four IBM 729 II, 729 IV, or 7330 Magnetic Tape Units. The program can modify itself according to information punched in a control card and thus performs a variety of sorting applications. This bulletin also provides information for preparing control cards and estimating the timing of sorting applications.

MINOR REVISION (January, 1963)

This publication is a reprint of an earlier edition with the format changed to conform to that of the Systems Reference Library.

Sort 1 for IBM 1401 — Specifications

Sort 1 is a generalized program designed to perform basic tape sorting functions for an IBM 1401 Data Processing System equipped with magnetic tape. It is classified as a generalized program because it is capable of modifying itself according to information punched in a control card by the user. This ability enables Sort 1 to perform a variety of sorting applications.

Sorting is taking data records that appear in some order on one or more reels of magnetic tape, rearranging these records in a particular sequence specified by the user, and rewriting them sequentially on tape.

General Information

A generalized tape sorting program such as Sort 1 has numerous commercial applications. For example, a wholesaler's daily transactions can be recorded as they occur. At the end of each day, Sort 1 can be utilized to write these transactions on tape in item number sequence, thus providing a compact and convenient daily business record.

Sort 1 performs applications such as this in two steps. The first, called Phase 1, is an *internal sort*. The records in random order are written in a semblance of their final sequence on two separate tape reels. Phase 2 is a *two-way merge*. This operation writes a single sequential tape file from the two reels that resulted from the internal sort in phase 1.

The Sort 1 program:

- sorts blocked or unblocked fixed-length records with a maximum block length of up to 800 characters
- sorts either numerical or alphanumerical records
- sorts according to control data contained in up to five fields of each record
- labels output tapes, if desired, in accordance with control card instructions

- provides a checkpoint routine that periodically writes the entire contents of core storage on tape and enables the user to stop and restart the program automatically at various stages of the program
- accommodates as many records as will fit on one reel of magnetic tape as the final output (input records may be contained on up to 99 reels)
- prints out blocks containing unreadable records or writes these blocks on a fifth tape unit, if available, called a *dump* tape (if a dump tape is unavailable, these blocks can be punched into cards).

Minimum Machine Requirements

Machine configuration requirements for IBM 1401 Sort 1 are minimal. The following features must be available:

- 4,000-character core storage capacity
- Minimum of four IBM 729 Model II, 729 Model IV, or 7330 Magnetic Tape Units (a fifth unit, if available, may be used as a dump tape)
- IBM 1402 Card Read-Punch
- IBM 1403 Printer
- High-Low-Equal Compare Feature

Sorting Technique

The sorting technique used in the Sort 1 program consists of reading a number of records from the input file, arranging them in short sequences, and writing these short sequences on alternate tapes. Subsequent passes merge these short sequences into longer sequences. By repeating this merging process, Sort 1 produces one long sequence called a sorted file.

One or two tape units are used for input, and two units are used for output during the initial sorting process. Two input units are used if the records to be sorted are contained on more than one reel. If the records are contained on more than two reels, the input units are alternated. The program automatically causes

a tape unit to stop and rewind when an end-of-file is reached, thus allowing the operator to change reels during input. In the merging process, the input units become output units, and vice versa.

For a more comprehensive discussion of tape sorting methods, refer to the IBM General Information Manual, *Sorting Methods for IBM Data Processing Systems*, Form F28-8001.

Sort 1 accomplishes the sorting operation in two steps – phase 1 and phase 2.

PHASE 1

1. Phase 1 writes the entire contents of core storage, after initialization, as the first record of the first output tape. This is known as the checkpoint procedure.
2. It reads into storage a number of input records, unblocks them if they are blocked, and sorts them internally.
3. It writes the short sequences on alternate output tapes.

PHASE 2

1. Phase 2 writes the entire contents of core storage as the first record of what has become the first output tape.
2. It merges the sequences written during phase 1 using as many passes as are required.
3. It reblocks the records according to the user's specifications and writes them as a sequential file on a single tape reel.

Allowable Input Record Configurations

Sort 1 accommodates only fixed-length input records. They may appear on tape either singly or in blocks. If input records are blocked, the number of records per block (blocking factor) must be constant for each job.

The blocking factor must be established so that a block contains no more than 800 characters or 730 characters, depending upon whether one, or more than one, control data fields are used in a particular job. If only one control data field is used, a block may contain up to 800 characters. For example, if each input record is 100 characters in length, the maximum blocking factor is 8 ($8 \times 100 = 800$). If more than one control data field is used, maximum block size is 730 characters. Maximum length for unblocked records is either 800 or 730 characters.

Input Blocking

Maximum input block size is determined by the number of positions of core storage set aside by the program for internal sorting during phase 1. If only one control data field is used, more storage area is freed

for other use, and 800 core locations are available for internal sorting. If more than one control data field is used, more storage positions are required by the program, and the area available for internal sorting is reduced to 730 positions. Thus, no more than either 800 or 730 characters at a time can be processed.

Processing time can be substantially reduced if an input block contains as close to the maximum as possible. For example, if maximum block size is 800, and a block contains 800 characters, only one READ operation is performed by the program before each internal sort. If a block contains 400 characters, two READS are performed by the program before the internal sort. If the block contains 200 characters, four READS are required, and so forth.

As the preceding case suggests, if the maximum allowable blocking factor is not used, a submultiple of it should be used. For example, assuming a maximum

RECORD LENGTH	MAXIMUM BLOCKING FACTOR	OTHER RECOMMENDED BLOCKING FACTORS	G
010-020	40	20, 10, 5, 4, 2, 1	40
021	38	19, 2, 1	38
022	36	18, 12, 9, 6, 4, 3, 2, 1	36
023	34	17, 2, 1	34
024	33	11, 3, 1	33
025	32	16, 8, 4, 2, 1	32
026	30	15, 10, 6, 5, 3, 2, 1	30
027	29	1	29
028	28	14, 7, 4, 2, 1	28
029	27	9, 3, 1	27
030	26	13, 2, 1	26
031-032	25	5, 1	25
033	24	12, 8, 6, 4, 3, 2, 1	24
034	23	1	23
035-036	22	11, 2, 1	22
037-038	21	7, 3, 1	21
039-040	20	10, 5, 4, 2, 1	20
041-042	19	1	19
043-044	18	9, 6, 3, 2, 1	18
045-047	17	1	17
048-050	16	8, 4, 2, 1	16
051-053	15	5, 3, 1	15
054-057	14	7, 2, 1	14
058-061	13	1	13
062-066	12	6, 4, 3, 2, 1	12
067-072	11	1	11
073-080	10	5, 2, 1	10
081-088	9	3, 1	9
089-100	8	4, 2, 1	8
101-114	7	1	7
115-133	6	3, 2, 1	6
134-160	5	1	5
161-200	4	2, 1	4
201-266	3	1	3
267-400	2	1	2
401-800	1	—	1

Figure 1. Recommended Blocking with One Control Data Field

allowable block size of 800 characters and an input record length of 50 characters, the maximum permissible blocking factor is 16 ($16 \times 50 = 800$). The blocking factor should be either 16 or a submultiple of 16 (8, 4, 2, 1). Blocking factors other than a submultiple (15, 14, 13, 12, 11, 10, 9, 7, 6, 5, 3) may be used, but they will cause an increase in total processing time.

Figure 1 contains maximum allowable blocking factors and other recommended blocking factors for all size records up to the maximum 800 characters. Figure 2 contains maximum allowable blocking factors and other recommended blocking factors for all size records up to the maximum 730 characters.

Output Blocking

Records may be blocked on the output tape according to the user's specifications, as punched in the control card. The output blocking factor must be such that

RECORD LENGTH	MAXIMUM BLOCKING FACTOR	OTHER RECOMMENDED BLOCKING FACTORS	G
010-018	40	20, 18, 8, 5, 4, 2, 1	40
019	38	19, 2, 1	38
020	36	18, 12, 9, 6, 4, 3, 2, 1	36
021	34	17, 2, 1	34
022	33	11, 3, 1	33
023	31	1	31
024	30	15, 10, 6, 5, 3, 2, 1	30
025	29	1	29
026	28	14, 7, 4, 2, 1	28
027	27	9, 3, 1	27
028	26	13, 2, 1	26
029	25	5, 1	25
030	24	12, 8, 6, 4, 3, 2, 1	24
031	23	1	23
032-033	22	11, 2, 1	22
034	21	7, 3, 1	21
035-036	20	10, 5, 4, 2, 1	20
037-038	19	1	19
039-040	18	9, 6, 3, 2, 1	18
041-042	17	1	17
043-045	16	8, 4, 2, 1	16
046-048	15	5, 3, 1	15
049-052	14	7, 2, 1	14
053-056	13	1	13
057-060	12	6, 4, 3, 2, 1	12
061-066	11	1	11
067-073	10	5, 2, 1	10
074-081	9	3, 1	9
082-091	8	4, 2, 1	8
092-104	7	1	7
105-121	6	3, 2, 1	6
122-146	5	1	5
147-182	4	2, 1	4
183-243	3	1	3
244-365	2	1	2
366-730	1	—	1

Figure 2. Recommended Blocking with More Than One Control Data Field

output block length does not exceed either 800 or 730 characters, depending upon the number of control data fields used. Maximum permissible input and output blocking factors are always the same for a particular job (see Figures 1 and 2). Note that processing time is reduced if the maximum permissible output blocking factor is used.

Maximum File Length

The input file to be processed by Sort 1 must be no longer than the number of records that can be contained on a single tape reel. This number will depend on record length, input blocking factor, and on whether processing is performed in the high- or low-density magnetic tape mode. The following formula enables the user to compute the maximum number of records that can be sorted in one job:

$$\text{Maximum Number of Input Records} = \frac{(K \times G)}{(G \times L) + \text{IRG}}$$

Explanation of symbols:

K = Number of character locations per tape reel

High-density tape — 15,350,000

Low-density tape — 5,520,000

IRG = Number of character locations per inter-record gap

High-density tape — 417

Low-density tape — 150

L = Characters per record

G = Largest multiple of input blocking factor that is less than or equal to either

$\frac{800}{L}$ when one control data field is used, or

$\frac{730}{L}$ when more than one control data field is used.

(See Figures 1 and 2 for values of G)

EXAMPLE

Compute the maximum file size for records 50 characters long with an input blocking factor of eight. Processing will be in the high-density mode and one control data field will be used. Referring to the preceding formula, the symbols will have the following value:

K = 15,350,000

IRG = 417

L = 50

G = 16

The formula is then evaluated as follows:

$$\frac{(15,350,000 \times 16)}{(16 \times 50) + 417} = 201,807$$

The maximum file size for this job is 201,807 records.

Tape Density

The Sort 1 program accommodates input reels written in either high- or low-density format; the final output reel may be written in either density, although high-density is recommended. The user need only set the density switch of the output tape unit to the desired

density. The tapes used for processing must be consistent in density, but they need not have the same density as the final output reel.

Note: If processing is performed in the high-density mode and final output is in low density, it is conceivable that the final output may require slightly more than one full reel of magnetic tape. In this situation, the program halts when an end-of-reel is encountered during final output. The user may then mount a new tape and press `START` to continue output.

Control Data Fields

From one to five fields of each input record can be specified to control sequencing. These fields can be located anywhere within the record, provided they are in the same place in each record. They can be of any length.

The location of each control field is specified by the user in the control card. If more than one control field is used, the user must specify which is to be compared first, which second, and so forth.

Although up to five fields are permitted, it is to the user's advantage to limit the number of control data fields to one. As noted previously, the use of only one control field raises the maximum permissible block size to 800. Secondly, processing time is reduced if the number of control fields is reduced. If more than one control field must be used, it is beneficial if the fields appear in the record sequentially in order of importance from left to right. Several fields can thus be treated by the program as one field. Control fields can contain any alphanumerical characters or special symbols. Standard collating sequence for the IBM 1401 is used.

Unreadable Input Records

Input tape blocks containing unreadable records (records that cause redundancy indications on one or more characters after several attempts at re-reading) may be treated in a variety of ways according to punches in the control card prepared by the user.

When an unreadable record is reached, the block containing it is read into storage, and the machine internally corrects the parity of the invalid character by either adding or removing the check bit. Thus, although the character is now valid for machine purposes, it may not be the same character that appeared on tape.

A punch in column 14 of the control card determines the next action taken on the block containing the unreadable record. The record can be corrected from the console, or the block containing it can be punched

into cards or written on a fifth tape unit, if available. If the unreadable record is corrected, the entire block will also be printed.

Unreadable records are corrected in the following manner. The program stops after the block containing the unreadable record is printed. This gives the user an opportunity to study the contents of the record. The user then has the option of continuing the sorting process with the record as it appears, or of correcting the invalid character manually before resuming processing. To continue sorting with the record as it appears, the user need only press the start key. To correct the invalid character, the user should:

1. turn on sense switch `G` and set the tape select switch to `D`
2. press `START`, causing the block containing the incorrect record to be re-read; the program again halts if the redundancy has not been corrected
3. manually load the correct character in its appropriate storage location
4. set the tape select switch back to `N`, and turn off sense switch `G`
5. press `START` to resume processing, beginning with the block that was just corrected.

Checkpoint and Restart

Because sorting is, by computer standards, a fairly lengthy procedure, a feature has been incorporated in the Sort 1 program that enables the user to stop processing at any stage of the sort if he must relinquish the machine. This same feature allows him to resume processing at a point in the program very close to where he stopped, thus saving considerable duplication of operating time.

The program accomplishes this by writing *checkpoints* periodically during the running of the sort. A checkpoint is a tape record containing the entire contents of storage. It is written as the initial record of the first output tape.

The first checkpoint is written during phase 1 after initialization and just before the reading of the first block of input records to be sorted. During phase 2 a checkpoint is written at the beginning of every merge pass.

If processing is stopped during phase 1, all sorting performed up to that point is lost, and the restart begins with the reading of the first block of records to be sorted. If processing is stopped during phase 2, only the merge pass that is interrupted is lost. The output of all preceding merge passes remains intact. When the program is interrupted, the user must, of course, save the output reels from the last pass and the reel containing the checkpoint.

To restart after an interruption, it is necessary only to:

1. mount the input and output reels
2. set the indicator of the tape unit on which the first output reel is placed to 1
3. press the tape load key

This automatically causes the first record of tape unit 1 (the checkpoint record) to be read into storage, and causes a branch to location 001 for the first instruction. This instruction causes the program to begin either at the beginning of phase 1 or at the beginning of the interrupted merge pass, depending on which checkpoint is used. The restart routine also causes a new checkpoint to be written.

To insure that the user sets the tape unit indicators to the correct numbers, the program automatically causes the numbers of the units being used to be printed out when the tape load key is pressed during the restart routine. The number of the pass during which the program was interrupted is also printed out. If the interruption occurred during phase 1, 00 is printed out.

Padding

The term *padding* refers to records added to a file to be sorted when the number of records in the file is not a multiple of the maximum allowable input blocking factor. These additional records are generated internally by the Sort 1 program.

Sort 1 automatically adds padding records to an input file if, preparatory to reading into storage the last block of records during phase 1, it finds that there are insufficient records to fill the processing area. Padding records generated by Sort 1 are sorted and merged in the same manner as data input records. They must, therefore, be composed either entirely of nines or entirely of blanks. The user's choice must be punched in the control card. If they contain nines, they will be the last records in the sorted file. If they contain blanks, they will be the first records of the sorted file.

EXAMPLE

Here is a case in which padding is required. An input file contains 90 records, and the maximum permissible input blocking factor is 16. The first five internal sorts process 16 records at a time. Prior to the sixth and final internal sort, however, only ten records remain to be read. Because the processing area of storage must contain the same number of input records during each internal sort, six padding records are read into the area at this point. Sixteen records are now ready for processing and the program continues. As this example indicates, padding can occur only in the final internal sort of phase 1.

Tape Labels

Sort 1 accommodates header labels on input reels and writes a header label on the final output reel, if it is desired by the user. When input header labels are specified, the program assumes that the header label is the first record of the reel. No provision is made for trailer labels. If they appear on input reels they are ignored by the program. Also, if one input reel contains a header label, all input reels must contain header labels, although the labels do not have to be the same in size or content. Maximum allowable header label length, on either input or output reels, is 80 characters.

Control Card Preparation

The user provides control information that enables Sort 1 to modify itself so that it can perform a particular application. Control information is supplied to the program by means of a single control card prepared by the user and inserted in the program deck.

When the control card is prepared, leading zeros are punched in fields containing information. For example, the field specifying the number of input reels (columns 3-4) is punched 05 if there are five input reels. Unused fields are left blank.

Control card format is shown in Figure 3. An explanation of each control card field follows.

Tape Unit Specification (Columns 1-4, 12, 13)

Four IBM 729 Model II, 729 Model IV or 7330 Magnetic Tape Units are required by the Sort 1 program. Two are used for input and two for output. The user specifies the number of each unit and the total number of tape reels on which his file is contained.

Column 1 is punched with the number of the first input tape unit. *Column 2* is punched with the number of the second input tape unit. *Columns 3-4* are punched with the total number of tape reels in which the file is contained. *Column 12* is punched with the number of the first output tape unit from phase 1. The checkpoint immediately prior to phase 1 is written as the first record of the reel mounted on this tape unit. *Column 13* is punched with the number of the second output tape unit from phase 1.

Blocking Information (Columns 5-11)

Columns 5-7 are punched with the number of characters per record. Note that only fixed-length records are permitted.

Columns 8-9 are punched with the input blocking factor. *Columns 10-11* are punched with the output blocking factor. Recommended blocking factors for

COLUMN NO.	DESCRIPTION
1	No. of first input tape unit
2	No. of second input tape unit
3-4	No. of input reels
12	No. of first output tape unit
13	No. of second output tape unit
5-7	Input record length
8-9	Input blocking factor
10-11	Output blocking factor
14	Unreadable record option
15	Tape density indicator
16	Input label indicator
17	Output label option
18	Padding character
19	No. of control data fields
20-22	No. of control data field characters
23-25	Location in record of control data field 1 (high-order position)
26-28	No. of characters in control data field 1
29-31	Location in record of control data field 2 (high-order position)
32-34	No. of characters in control data field 2
35-37	Location in record of control data field 3 (high-order position)
38-40	No. of characters in control data field 3
41-43	Location in record of control data field 4 (high-order position)
44-46	No. of characters in control data field 4
47-49	Location in record of control data field 5 (high-order position)
50-52	No. of characters in control data field 5
53-79	Unused
80	Tape mark option

Figure 3. Sort 1 Control Card

records of various lengths are shown in Figures 1 and 2.

Unreadable Record Option (Column 14)

As noted previously, the action taken on blocks containing unreadable records is determined by control card specifications. The control card offers the option of punching the block into cards, writing it on a fifth tape unit, or allowing the user to correct the unreadable record manually. In the latter case, the entire block is also printed.

Column 14 is punched with the unreadable record option. If blocks containing unreadable records are to be punched into cards, column 14 is left blank. If blocks containing unreadable records are to be written on a dump tape, the number of this fifth unit is punched in column 14. If unreadable records are to be corrected from the console, a zero is punched in column 14.

Tape Density Indicator (Column 15)

The tapes used in processing may be written in either high- or low-density format, regardless of the density

of the input tapes. Column 15 is punched with the tape density indicator. If these tapes are to be low density, a zero is punched in column 15. If they are to be high density, a 1 is punched in column 15. The density switches of the tape units must be set to the appropriate density. The density of the final output tape need not be the same as the density of the processing tapes. High density is recommended for processing and final output.

Tape Labels (Columns 16, 17, and 80)

The user specifies in these columns the presence or absence of header labels on input reels and whether or not a header label is desired on the output reel. The Sort 1 program ignores trailer labels.

Column 16 contains the input label indicator. Column 16 is left blank if the input reels do not contain header labels. If the input reels contain header labels, a 1 is punched in column 16. Note that if a 1 is punched in column 16, every input reel must have a label as its first record. Each label may or may not be followed by a tape mark.

Column 17 contains the output label option. If the output reel is to have the same label as the first input reel, a 1 is punched in column 17. If the output reel is to have a new label, a 2 is punched in column 17. In the latter case a card punched with the contents of the label must be provided by the user and included with the program deck and the control card. If there is to be no label on the output reel, column 17 is left blank.

Column 80 contains the tape mark option for output labels. If the output reel is to have a header label, the user has the option of writing a tape mark immediately after it. If the output label is to be followed by a tape mark, a zero is punched in column 80. If no tape mark is desired, column 80 is left blank.

Padding (Column 18)

Column 18 is punched with the character to be used throughout the program in padding records. If nines are desired as padding records, column 18 must contain a nine. If blanks are desired as padding records, column 18 is left blank.

Control Data Specifications (Columns 19-52)

The Sort 1 program bases record sequence on the contents of up to five control data fields contained in each input record. These fields, specified by the user, are compared from record to record.

The control data fields, if there are more than one, do not have to be contiguous, nor do they have to appear in the record in the same order in which they will be compared. The user specifies control fields in the control card in the order of their importance. Thus,

the control field to be compared first is designated control data field 1, even though it may appear in the input record following control data field 2.

Column 19 is punched with the total number of control data fields used. Valid punches in this column are 1 through 5.

Columns 20-22 are punched with the total number of characters in the up-to-five control data fields used. This number is limited only by the size of the record.

Columns 23-28, 29-34, 35-40, 41-46, and 47-52 are punched with the specifications of control data fields 1, 2, 3, 4, and 5, respectively. The first three columns for each field (columns 23-25 for control data field 1) are punched with the location in the record of the high-order character of the control field. The first location of every input record is considered 001. The second three columns for each field (columns 26-28 for control data field 1) are punched with the total number of characters in the control field.

If less than five control data fields are used, unused control field columns in the control card must be left blank. For example, if the user specifies two control data fields for a particular job, columns 35-52 of the control card must be blank.

EXAMPLE

Here is an example of control data field specification. In an input file containing records 80 characters long, three control fields are used. The first (major) control field to be compared occupies locations 71-80. The second (intermediate) control field to be compared occupies locations 6-10. The third (minor) control field to be compared occupies locations 28-34. Figure 4 shows the punches required for this example in columns 19-52 of the control card.

Unused Columns (Columns 53-79)

Columns 53-79 of the control card are not used by the Sort 1 program and may be either punched for identification purposes or left blank, at the discretion of the user.

CARD COLUMNS	PUNCH	EXPLANATION
19	3	Total number of control data fields
20-22	022	Total number of characters in control data fields
23-25	071	High-order position of control data field 1
26-28	010	Number of characters in control data field 1
29-31	006	High-order position of control data field 2
32-34	05	Number of characters in control data field 2
35-37	028	High-order position of control data field 3
38-40	07	Number of characters in control data field 3
41-52	blank	Only three control data fields used

Figure 4. Example of Control Data Field Specification

R/G	p
1-2	1
3-4	2
5-8	3
9-16	4
17-32	5
33-64	6
65-128	7
129-256	8
257-512	9
513-1024	10
1025-2048	11
2049-4096	12
4097-8192	13
8193-16384	14
16385-32768	15

Figure 5. Number of Merge Passes

Estimating Sorting Times

To estimate the time required by the Sort 1 program to process a given file of records, it is necessary to know the duration of phase 1, the duration of each merge pass of phase 2, the number of merge passes required in phase 2, the tape time required for each pass, and the time required for tape rewinding.

All of these figures can be obtained from the tables in Figures 5, 6, and 7. Figure 5 lists the number of merge passes required for various job sizes. Figure 6 provides the remainder of the data required for solving the timing formula if only one control data field is used. Figure 7 provides the remainder of the data required for solving the timing formula if more than one control field is used.

Number of Merge Passes

The table in Figure 5 contains the number of merge passes (p) required for various jobs. In order to determine the number of passes, it is necessary to divide the number of records in the job (R) by the maximum permissible blocking factor (G). Various values for G are shown in Figures 1 and 2.

For example, in a particular job there are 100,000 input records (R) of 80 characters each. Only one control data field is utilized, so G can equal 10 (see Figure 1). Using the formula in Figure 5, $R/G = 10,000$. The number of merge passes required is therefore 14.

It should be noted at this point that during phase 2 the program takes advantage of any sequencing already existing in the user's file. If a degree of sequencing is present, the number of merge passes in phase 2 is reduced. Experience has shown that pre-existing sequencing in an unsorted file may reduce

the number of merge passes (p) by an average of 1/7. Thus the actual number of merge passes required in the example in the previous paragraph is more likely to be 12. The reduction in number of merge passes depends upon the degree of existing sequencing in a file. The user should take this factor into consideration when calculating the value of p.

Timing Formula

The following timing formula requires the use of many factors, the values for which can be determined from the tables in Figures 5, 6, and 7. This formula provides an accurate timing estimate if the file to be sorted has the following characteristics:

1. Maximum permissible blocking factor is used for input and output.
2. Maximum tape rewind time is required for each merge pass.
3. Only one control data field is used.

If any of these conditions are not met, the factors in the basic timing formula must be adjusted to make the formula accurate. These adjustments are described subsequently. Here is the basic timing formula:

$$\text{Total Time (Minutes)} = \frac{P1 \times R}{60,000} + \frac{P2 \times R(p)}{60,000} + \frac{(T)(R)(p+1)}{60,000} + W(p+1)$$

Where:

P1 = Process time for phase 1 in milliseconds per record (see Figure 6)

P2 = Process time for each pass of phase 2 in milliseconds per record (see Figure 6)

R = Total number of records to be sorted

p = Number of merge passes in phase 2 (see Figure 5)

T = Tape time for phase 1 and each pass of phase 2 in milliseconds per record*

T2L = IBM 729 Model II Low Density (see Figure 6)

T2H = IBM 729 Model II High Density (see Figure 6)

T4L = IBM 729 Model IV Low Density (see Figure 6)

T4H = IBM 729 Model IV High Density (see Figure 6)

W = Rewind time

IBM 729 Model II = 1.2 minutes

IBM 729 Model IV = 0.9 minutes

As noted previously, this timing formula is accurate only if all of the foregoing conditions are met. If the input and/or output blocking factor are not the maximum permissible, processing time is increased. Similarly, if more than one control data field is utilized, processing time is increased.

EXAMPLE

This example illustrates the use of the Sort 1 timing formula without any adjustment necessary. Assume

that a file of records to be sorted has the following specifications:

Record length (L) = 30
 Number of records (R) = 100,000
 Input and output blocking factor = 26
 Maximum permissible blocking factor (G) = 26
 Number of control fields = 1
 Length of control field (CF) = 6

IBM 729 Model II Magnetic Tape Units in the low-density mode are to be used.

It is first necessary to consult the table in Figure 5 to determine how many merge passes (p) are required in phase 2 of this operation. Because $\frac{R}{G}$ in this case is $\frac{100,000}{26}$, or 3,846, p = 12.

The table in Figure 6 is then consulted to obtain the values of P1, P2, and T. (T in this example is T2L, because tape operations are performed on a Model II tape unit in the low-density mode.)

The timing chart in Figure 6 is read by first scanning the column labeled (L) to find the appropriate record length, in this case 30. Within this group, the column labeled (CF) is scanned to find the appropriate control field length, in this case 6. By reading across this line, the user finds that P1 = 15.3, P2 = 3.6, and T2L = 4.9. Because processing is on a Model II tape unit, rewind time (W) = 1.2.

The proper figures can now be inserted into the timing formula as follows:

$$\begin{aligned} \text{Total time (minutes)} &= \frac{15.3(100,000)}{60,000} + \frac{(3.6)(100,000)(12)}{60,000} + \\ &\quad \frac{(4.9)(100,000)(12+1)}{60,000} + 1.2(12+1) \\ &= 219.3 \text{ minutes} \end{aligned}$$

Rewind Time Considerations

The formula for total sort time includes the term W, the time required to rewind a full reel of magnetic tape. An IBM 729 Model II Magnetic Tape Unit requires 1.2 minutes to rewind a full reel. An IBM 729 Model IV Magnetic Tape Unit requires 0.9 minute to rewind a full reel. If more than 450 feet of tape must be rewound, total rewind time is not reduced substantially enough to affect total sorting time. If 450 feet or less are to be rewound, however, rewind time is considerably lessened. This smaller rewind time can be substituted in the timing formula for W, in order to give a more accurate total time estimate.

The file size corresponding to 450 feet of tape is determined by multiplying the value of the maximum file size for a particular job by 0.2. Maximum file size is indicated in Figures 6 and 7 in the columns labeled R_H (for high density) and R_L (for low density). If the total number of records to be sorted (R) is equal to

* Sort 1 timing information relating to 1401 systems equipped with IBM 7330 Magnetic Tape Units will be presented in a subsequent publication.

L	G	CF	P1	P2	T2H	T4H	T2L	T4L	R _H	R _L
10	40	2	19.4	2.98	1.02	.68	1.88	1.24	749,000	386,000
		4	19.9	3.02	1.02	.68	1.88	1.24	749,000	386,000
20	40	2	19.6	3.21	1.50	1.00	3.22	2.12	504,000	223,000
		4	20.1	3.26	1.50	1.00	3.22	2.12	504,000	223,000
30	26	2	14.6	3.5	2.3	1.5	4.9	3.2	333,000	148,500
		4	15.0	3.5	2.3	1.5	4.9	3.2	333,000	148,500
		6	15.3	3.6	2.3	1.5	4.9	3.2	333,000	148,500
40	20	2	12.6	3.7	3.0	2.0	6.4	4.3	252,000	111,700
		5	13.1	3.8	3.0	2.0	6.4	4.3	252,000	111,700
		10	13.7	3.9	3.0	2.0	6.4	4.3	252,000	111,700
60	13	5	10.9	4.3	4.5	3.0	9.7	6.4	166,000	74,000
		10	11.4	4.4	4.5	3.0	9.7	6.4	166,000	74,000
		15	11.9	4.6	4.5	3.0	9.7	6.4	166,000	74,000
80	10	5	10.3	4.8	6.0	4.0	12.9	8.5	126,000	55,800
		10	10.7	4.9	6.0	4.0	12.9	8.5	126,000	55,800
		15	11.1	5.1	6.0	4.0	12.9	8.5	126,000	55,800
		20	11.5	5.2	6.0	4.0	12.9	8.5	126,000	55,800
100	8	5	7.7	5.3	7.5	5.0	16.1	10.6	100,300	44,600
		10	10.4	5.5	7.5	5.0	16.1	10.6	100,300	44,600
		15	10.8	5.6	7.5	5.0	16.1	10.6	100,300	44,600
		20	11.2	5.7	7.5	5.0	16.1	10.6	100,300	44,600
120	6	5	10.0	5.9	9.4	6.3	20.0	13.0	80,800	36,500
		10	10.3	6.0	9.4	6.3	20.0	13.0	80,800	36,500
		15	10.6	6.1	9.4	6.3	20.0	13.0	80,800	36,500
		20	10.9	6.2	9.4	6.3	20.0	13.0	80,800	36,500
150	5	10	10.7	6.8	11.5	7.7	24.4	16.1	65,600	29,400
		15	11.0	6.9	11.5	7.7	24.4	16.1	65,600	29,400
		20	11.3	7.0	11.5	7.7	24.4	16.1	65,600	29,400
		25	11.6	7.1	11.5	7.7	24.4	16.1	65,600	29,400
200	4	10	11.8	8.0	15.0	10.1	32.2	21.3	50,400	22,300
		20	12.3	8.3	15.0	10.1	32.2	21.3	50,400	22,300
		25	12.5	8.4	15.0	10.1	32.2	21.3	50,400	22,300
		30	12.8	8.5	15.0	10.1	32.2	21.3	50,400	22,300
300	2	10	11.7	10.9	25.2	16.9	51.0	33.7	30,100	14,200
		20	12.0	11.1	25.2	16.9	51.0	33.7	30,100	14,200
		30	12.4	11.4	25.2	16.9	51.0	33.7	30,100	14,200
		35	12.5	11.5	25.2	16.9	51.0	33.7	30,100	14,200
400	2	15	14.0	13.3	30.0	20.1	64.4	42.5	25,200	11,150
		20	14.3	13.4	30.0	20.1	64.4	42.5	25,200	11,150
		30	14.7	13.6	30.0	20.1	64.4	42.5	25,200	11,150
		40	15.0	13.9	30.0	20.1	64.4	42.5	25,200	11,150
500	1	15	14.2	16.7	45.6	30.6	88.6	58.6	16,700	8,150
		25	14.5	17.0	45.6	30.6	88.6	58.6	16,700	8,150
		35	14.7	17.2	45.6	30.6	88.6	58.6	16,700	8,150
		45	14.9	17.4	45.6	30.6	88.6	58.6	16,700	8,150
600	1	20	16.6	19.1	50.4	33.8	102.0	67.4	15,000	7,070
		30	16.9	19.4	50.4	33.8	102.0	67.4	15,000	7,070
		40	17.1	19.6	50.4	33.8	102.0	67.4	15,000	7,070
		50	17.3	19.8	50.4	33.8	102.0	67.4	15,000	7,070
700	1	20	19.0	21.4	55.2	37.0	115	76.2	13,700	6,240
		30	19.2	21.6	55.2	37.0	115	76.2	13,700	6,240
		40	19.4	21.9	55.2	37.0	115	76.2	13,700	6,240
		50	19.6	22.2	55.2	37.0	115	76.2	13,700	6,240
800	1	20	21.2	23.7	60.0	40.2	129	85.0	12,600	5,580
		30	21.5	24.0	60.0	40.2	129	85.0	12,600	5,580
		40	21.7	24.2	60.0	40.2	129	85.0	12,600	5,580
		50	21.9	24.4	60.0	40.2	129	85.0	12,600	5,580

Figure 6. Timing Factors for Files with One Control Data Field

or less than $R_H \times 0.2$ (high density) or $R_L \times 0.2$ (low density), then W is reduced in value when used in the timing formula.

Furthermore, the user should note that during phase 2, the reels being rewind contain, in general, only half the entire file. This is an approximation, and therefore it is wise to be conservative and assume that on each rewind $\frac{3}{4}$ of the full file is being rewind. This further reduces the value of W . The following formulas are used to determine the value of W :

Model II High Density:

$$\text{Rewind Time (W)} = \frac{.75R}{R_H \times 0.2} \times 1.2$$

Model II Low Density:

$$\text{Rewind Time (W)} = \frac{.75R}{R_L \times 0.2} \times 1.2$$

Model IV High Density:

$$\text{Rewind Time (W)} = \frac{.75R}{R_H \times 0.2} \times 0.9$$

Model IV Low Density:

$$\text{Rewind Time (W)} = \frac{.75R}{R_L \times 0.2} \times 0.9$$

An example of how this reduced value of W could occur can be seen by referring to the previous timing example. Assume that the number of records in the file (R), is not 100,000, but 10,000. The table in Figure 6 shows that the maximum file size for this job (R_L) is 148,500. This figure multiplied by 0.2 is 29,700, so it is clear that a file containing only 10,000 records occupies less than 450 feet of magnetic tape. One of the formulas can thus be used to reduce the value of W . Because processing is performed on a Model II tape unit in the low-density mode, the applicable formula is:

$$W = \frac{.75R}{R_L \times 0.2} \times 1.2. \text{ Thus:}$$

$$W = \frac{(.75)(10,000)}{(148,500)(0.2)} \times 1.2, \text{ or}$$

$$W = 0.3 \text{ (approx.)}$$

This new value for W is inserted in the basic timing formula, replacing 1.2.

Control Data Field Considerations

If more than one field per record is used to control sequencing in a particular file, the factors used in the basic timing formula must be adjusted to reflect the increase in processing time. Such considerations as the number of control data fields used, and the length of the control data fields, affect total sorting time.

When more than one control data field is used, the values for $P1$, $P2$, T , R_H , and R_L should be determined from the table in Figure 7. These will be the values used in the basic timing formula.

Before the basic timing formula can be used, however, the values of $P1$ and $P2$ must be adjusted to reflect the number and length of additional control data fields. The following formulas are used to compute the

values of $\Delta P1$ and $\Delta P2$. These values must be added to the values of $P1$ and $P2$, as indicated in the table in Figure 7, before solving the basic timing formula.

$$\Delta P1 = \left(\frac{G+3}{4} + \frac{1}{G} \right) [2LAF + 153(NAF - 1) + 114] .0115$$

(Note: The formula just given for $\Delta P1$ is valid only if the value of G is 4 or greater. When $G = 1, 2,$ or 3 , one of the following formulas should be used for $\Delta P1$.)

If $G = 1$,

$$\Delta P1 = .0115 [2LAF + 153(NAF - 1) + 114]$$

If $G = 2$,

$$\Delta P1 = 1.5 [2LAF + 153(NAF - 1) + 114] .0115$$

If $G = 3$,

$$\Delta P1 = 1.7 [2LAF + 153(NAF - 1) + 114] .0115$$

The following formula is used at all times to determine the value of $\Delta P2$:

$$\Delta P2 = \left(\frac{G+1}{G} \right) [2LAF + 153(NAF - 1) + 114] .0115$$

In all of the foregoing formulas,

LAF = Length of *additional* fields

NAF = Number of *additional* fields

EXAMPLE

This example illustrates the use of the Sort 1 timing formula when the file to be sorted has more than one control data field per record. Assume that the job has the following specifications:

Record length (L) = 80

Number of records (R) = 60,000

Input and output blocking factor = 9

Maximum permissible blocking factor (G) = 9

Number of control data fields = 3

Length of first control data field (CF) = 5

Length of second control data field = 6

Length of third control data field = 4

IBM 729 Model IV Magnetic Tape Units in the high-density mode are used.

The table in Figure 5 indicates that the number of merge passes (p) required in phase 2 is 13, because

$$\frac{R}{G} = \frac{60,000}{9}, \text{ or } 6,667.$$

Because processing is performed on a Model IV tape unit in the high-density mode, T will assume the value of $T4H$, and the rewind time (W) = 0.9. Because this job uses more than one control data field, the table in Figure 7 must be referred to for the values to be used in the timing formula. In this table, the length of the first control data field to be compared is the number to be searched for in the column labeled CF. The table in Figure 7 indicates that for this particular job, $P1 = 10.0$, $P2 = 4.6$, and $T4H = 4.2$. It must be remembered, however, that the values of $P1$ and $P2$ must be incremented by the value of $\Delta P1$ and $\Delta P2$ respectively before the timing formula can be solved.

L	G	CF	P1	P2	T2H	T4H	T2L	T4L	R _H	R _L
20	36	2	18.1	3.20	1.56	1.05	3.28	2.16	484,000	220,000
		4	18.5	3.26	1.56	1.05	3.28	2.16	484,000	220,000
30	24	2	13.9	3.5	2.3	1.6	4.9	3.3	323,000	146,000
		4	14.2	3.5	2.3	1.6	4.9	3.3	323,000	146,000
		6	14.5	3.6	2.3	1.6	4.9	3.3	323,000	146,000
40	18	2	11.9	3.6	3.1	2.1	6.6	4.3	242,000	109,700
		5	12.3	3.7	3.1	2.1	6.6	4.3	242,000	109,700
		10	12.9	3.8	3.1	2.1	6.6	4.3	242,000	109,700
60	12	5	10.5	4.2	4.7	3.1	9.8	6.5	161,500	73,200
		10	11.0	4.3	4.7	3.1	9.8	6.5	161,500	73,200
		15	11.5	4.4	4.7	3.1	9.8	6.5	161,500	73,200
80	9	5	10.0	4.6	6.2	4.2	13.1	8.7	121,000	54,900
		10	10.3	4.7	6.2	4.2	13.1	8.7	121,000	54,900
		15	10.7	4.9	6.2	4.2	13.1	8.7	121,000	54,900
		20	11.1	5.0	6.2	4.2	13.1	8.7	121,000	54,900
100	7	5	9.8	5.4	7.9	5.3	16.5	10.9	95,800	43,700
		10	10.1	5.5	7.9	5.3	16.5	10.9	95,800	43,700
		15	10.4	5.6	7.9	5.3	16.5	10.9	95,800	43,700
		20	10.8	5.7	7.9	5.3	16.5	10.9	95,800	43,700
120	6	5	10.0	5.9	9.4	6.3	20.0	13.0	80,800	36,500
		10	10.3	6.0	9.4	6.3	20.0	13.0	80,800	36,500
		15	10.6	6.1	9.4	6.3	20.0	13.0	80,800	36,500
		20	10.9	6.2	9.4	6.3	20.0	13.0	80,800	36,500
150	4	10	10.6	6.9	12.6	8.4	25.6	16.8	60,300	28,300
		15	10.9	7.0	12.6	8.4	25.6	16.8	60,300	28,300
		20	11.1	7.1	12.6	8.4	25.6	16.8	60,300	28,300
		25	11.4	7.2	12.6	8.4	25.6	16.8	60,300	28,300
200	3	10	10.4	8.2	16.8	11.3	34.0	22.5	45,100	21,200
		20	10.9	8.5	16.8	11.3	34.0	22.5	45,100	21,200
		25	11.1	8.6	16.8	11.3	34.0	22.5	45,100	21,200
		30	11.4	8.7	16.8	11.3	34.0	22.5	45,100	21,200
300	2	10	11.7	10.9	25.2	16.9	51.0	33.7	30,100	14,200
		20	12.0	11.1	25.2	16.9	51.0	33.7	30,100	14,200
		30	12.4	11.4	25.2	16.9	51.0	33.7	30,100	14,200
		35	12.5	11.5	25.2	16.9	51.0	33.7	30,100	14,200
400	1	15	11.9	14.4	40.8	27.4	75.2	49.8	18,700	9,640
		20	12.0	14.5	40.8	27.4	75.2	49.8	18,700	9,640
		30	12.3	14.8	40.8	27.4	75.2	49.8	18,700	9,640
		40	12.5	15.0	40.8	27.4	75.2	49.8	18,700	9,640
500	1	15	14.2	16.7	45.6	30.6	88.6	58.6	16,700	8,150
		25	14.5	17.0	45.6	30.6	88.6	58.6	16,700	8,150
		35	14.7	17.2	45.6	30.6	88.6	58.6	16,700	8,150
		45	14.9	17.4	45.6	30.6	88.6	58.6	16,700	8,150
600	1	20	16.6	19.1	50.4	33.8	102	67.4	15,000	7,070
		30	16.9	19.4	50.4	33.8	102	67.4	15,000	7,070
		40	17.1	19.6	50.4	33.8	102	67.4	15,000	7,070
		50	17.3	19.8	50.4	33.8	102	67.4	15,000	7,070
700	1	20	19.0	21.4	55.2	37.0	115	76.2	13,700	6,240
		30	19.2	21.6	55.2	37.0	115	76.2	13,700	6,240
		40	19.4	21.9	55.2	37.0	115	76.2	13,700	6,240
		50	19.6	22.2	55.2	37.0	115	76.2	13,700	6,240
730	1	20	19.6	22.2	56.6	38.0	119	78.8	13,350	6,030
		30	19.9	22.4	56.6	38.0	119	78.8	13,350	6,030
		40	20.1	22.6	56.6	38.0	119	78.8	13,350	6,030
		50	20.3	22.8	56.6	38.0	119	78.8	13,350	6,030

Figure 7. Timing Factors for Files with More Than One Control Data Field

In solving the formulas for $\Delta P1$ and $\Delta P2$, the symbols have the following value:

$G = 9$
 $LAF = 10$
 $NAF = 2$

The formulas are then solved as follows:

$$\Delta P1 = \left(\frac{9+3}{4} + \frac{1}{9} \right) [2(10) + 153(2-1) + 114] .0115$$

$$= 10.3$$

$$\Delta P2 = \left(\frac{9+1}{9} \right) [2(10) + 153(2-1) + 114] .0115$$

$$= 3.7$$

These values are then added to P1 and P2 so that $P1 = 20.3$ and $P2 = 8.3$. The basic timing formula can then be solved as follows:

$$\text{Total Time (minutes)} = \frac{20.3(60,000)}{60,000} + \frac{(8.3)(60,000)(13)}{60,000} + \frac{(4.2)(60,000)(14)}{60,000} + 0.9(14)$$

$$= 199.7 \text{ minutes}$$

It is important to note that although three control data fields are present in each record, the program will not necessarily have to compare all three fields from each record to establish sequence. For example, it may be necessary to compare the second control fields in only 1/20th of the records, and it may be necessary to compare the third control fields in only 1/50th of the records. The values of $\Delta P1$ and $\Delta P2$ can thus be correspondingly decreased to about 1/35th of their respective sizes to more accurately reflect the time required. Whether an adjustment of this type need be made, and the amount of the adjustment, must be determined by the user.

Blocking Considerations

In all the sorting situations covered so far in this discussion of timing estimates, the input and output blocking factors have been equal to G. That is, both have been the maximum permissible. If either the input blocking factor (B_i) or the output blocking factor (B_o) is less than G, however, several adjustments must be made to the basic timing formula.

If B_i is less than G, timing for phase 1 must be increased. If B_o is less than G, timing for the *last* pass of phase 2 must be increased. In order to facilitate these adjustments, the basic timing formula has been rearranged, or sectionalized, to show the timing for phase 1, the last merge pass of phase 2, all other merge passes, and rewind time. This sectionalized timing formula is shown in Figure 8.

IF INPUT BLOCKING FACTOR IS LESS THAN G...

In this case the value of P1 and the value of T must be increased in the portion of the timing formula that indicates the timing for phase 1 (see Figure 8). The following formulae indicate the amounts by which P1 and T are incremented:

$$\Delta P1 = \frac{1.4(G-B_i)}{G(B_i)} + \frac{0.9}{G}$$

$$\left. \begin{array}{l} \Delta T2H \\ \Delta T2L \end{array} \right\} = 10.8 \left(\frac{G-B_i}{G(B_i)} \right)$$

$$\left. \begin{array}{l} \Delta T4H \\ \Delta T4L \end{array} \right\} = 7.3 \left(\frac{G-B_i}{G(B_i)} \right)$$

IF OUTPUT BLOCKING FACTOR IS LESS THAN G...

In this case the value of P2 and the value of T must be increased in the portion of the timing formula that indicates the timing for the *last* merge pass of phase 2 (see Figure 8). The following formulae indicate the amounts by which P2 and T are incremented:

$$\Delta P2 = \frac{1.2(G-B_o)}{G(B_o)}$$

$$\left. \begin{array}{l} \Delta T2H \\ \Delta T2L \end{array} \right\} = 10.8 \left(\frac{G-B_o}{G(B_o)} \right)$$

$$\left. \begin{array}{l} \Delta T4H \\ \Delta T4L \end{array} \right\} = 7.3 \left(\frac{G-B_o}{G(B_o)} \right)$$

Although total processing time is increased when either input or output blocking is less than G, it is important to note that the degree of increase depends upon the size of the file being sorted. In lengthy files, the difference in sorting time is almost insignificant. As files become progressively shorter, however, the percentage of increase becomes more substantial.

	Phase 1	Last pass of phase 2	All other merge passes	Rewind time
Total time (minutes)	$\frac{P1(R)}{60,000} + \frac{T(R)}{60,000}$	$\frac{P2(R)}{60,000} + \frac{T(R)}{60,000}$	$\frac{(P2)(R)(p-1)}{60,000} + \frac{(T)(R)(p-1)}{60,000}$	$W(p+1)$
	$\underbrace{\hspace{10em}}_{\text{Must be adjusted if } B_i \neq G}$	$\underbrace{\hspace{10em}}_{\text{Must be adjusted if } B_o \neq G}$		

Figure 8. Sectionalized Version of Basic Sort 1 Timing Formula

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, with a small registered trademark symbol (®) to the right.

Technical Newsletter

File No. 1401-33

Re: Form No. J24-1422-1

This Newsletter No. N24-0162

Date: December 23, 1963

Previous Newsletter Nos. None

A change to the publication: Sort 1 for IBM 1401: Specifications, J24-1422-1.

Page 3, Output Blocking. Add:

If the output blocking factor is not a submultiple of G, the following message is printed:

BO NOT FACTOR OF B * SET BO TO B

The sort program replaces BO (output blocking factor) with B (maximum blocking factor) and continues processing.



International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, New York