



IBM

Reference Manual

IBM 7070/7074 Generalized Sorting Program

Sort 90

IBM[®]

Reference Manual
IBM 7070/7074 Generalized Sorting Program
Sort 90

CONTENTS

	<u>Page</u>
Introduction	1
Program Description	4
Phase I	4
Phase II	5
Phase III	5
Example	5
Tape Allocation and Handling	8
Input Tape Units	8
Input Tape Units vs. Work Tape Units	10
Output Tape Units	10
Effects of Number of Tape Channels	10
Miscellaneous Tape Units	11
Input File Considerations	12
Record Types	12
Record and Block Sizes	12
Number of Records Allowable	13
Number of Reels Allowable	13
Control Data Specifications	15
Control Data Fields and Segments	15
Optimal Control Data Arrangement	16
Collating Sequences	17
Miscellaneous Features and Provisions	18
Record Counts	18
Hash Totals	18
Phase I Area Allocation	19
Checkpoint and Restart	20
Unreadable Records Procedure	20
Sorting Tape Files Prepared by other IBM Data Processing Systems	21
Tape Labels	21
Operation	23
Initiation of Sort 90	23
Actions During Sort 90	23
Interruption and Restart	24
Console Switches and Dials	25

Control Card Preparation	27
Control Card Reading	27
Control Card Checking	27
Control Card Listing	27
Messages	34
Phase I	34
Phase II	35
Phase III	36
Programmed Halts	37
Phase I	37
Phase II	45
Phase III	47
Modifications	50
Patching	50
Types of Modifications	50
Placement of Modifications	51
Settings of Internal Latches	52
Checkpoint and Restart	52
Editing	52
Summarizing	53
Phase I Modifications	55
Phase II Modifications	62
Phase III Modifications	66
Appendix I: Glossary of Tape Sorting and Tape Handling Terms . .	74
Appendix II: Calculation of Sorting Times	76
Appendix III: Non-ending Sort Condition	85
Index	86

INTRODUCTION

IBM 7070 Sort 90 is a generalized tape sorting program written for use on an IBM 7070 Data Processing System equipped with IBM 729 II and/or IV Magnetic Tape Units. Sort 90 is considered a "generalized" sort program because it is capable of modifying itself according to the specifications detailed on control cards. It can thus perform a variety of sorting applications and will adapt itself to perform the specific job desired by the user. Since in any one 7070 installation there may be a great number of tape sorting applications, each one differing from the others, the ability to use the 7070 Sort 90 program for each of these applications leads to obvious economies in programming time and operator effort.

The purpose of the 7070 Sort 90 program is to cause the ordering of previously unordered records of a tape file. A companion program, 7070 Merge 91, can combine as many as eight sorted files into a single ordered sequence. Many combined usages of the two programs will suggest themselves. For example, Sort 90 might be used on a daily basis to sort the records of a transaction file. At the end of the week, Merge 91 could be used to merge together the five daily sorted transaction files in preparation for a file maintenance operation involving a master file and the weekly transaction file.

Sort 90 is a large program and is composed of three phases: Phase I, which performs a preliminary sequencing of records of the file to be sorted; Phase II, which through repeated computer passes brings the records into very nearly completely sorted order; and Phase III, which completes the sorting process and writes a sorted output file.

Some of the features of 7070 Sort 90 are:

Large numbers of records may be sorted in one application of Sort 90; the maximum possible number of records for a sort is related to the number of tape units which are available to the program.

Records expressed in any one of three different configurations may be sorted. The configurations are those designated by the 7070 Input/Output Control System as Forms 1, 2, and 3, respectively. Extremely wide latitudes are permitted as to record length and blocking factor or block size. Output blocks may be different in size from input blocks.

The control data for sorting may consist of as many as ten control data fields, each located in any positions of the record. The control data may include as many as 160 digits of information.

Any one of four different collating sequences may be specified by the user for control over the order into which Sort 90 is to sort records.

Complete advantage is taken of any "natural" sequencing which occurs in the sort input file. The Sort 90 running time will be shorter when a partially ordered file is being sorted than when input records are in random order.

A number of options are included in Sort 90 as to the accumulating of record counts and/or hash totals, so as to yield additional controls over record handling beyond the standard 7070 Input/Output Control System block count.

Sort 90 will operate on an IBM 7070 with either 5,000 or 10,000 words of core storage. In either case, the exact amount of storage to be used by the program is controllable by the user. Sort 90 will benefit in speed from being allowed as much storage as is available.

The Sort 90 program may be loaded from cards or from tape. The control information it requires may be read through an IBM 7500 Card Reader, an IBM 7501 Console Card Reader, or a tape unit, or may alternately be left in 7070 storage by a program preceding Sort 90.

The magnetic tape units to be used in any application of Sort 90 are specified by, and are completely controllable by, the user.

Sort 90 will perform reading, checking, preparation, and writing of header and trailer tape labels on all input tapes, work tapes, and output tapes.

Sort 90 may be interrupted at the beginning of, or during any pass, except Phase I, and later restarted from the beginning of that pass.

It is recognized that for some applications the user may wish to alter some features of the operation of Sort 90, or to add to the program other functions not specifically a part of sorting, such as editing and summarizing of records. Accordingly, an extensive framework is provided for the swift and convenient addition to Sort 90 of modifications. A separate section of this manual is devoted to a discussion of the specific modification possibilities.

This manual is to serve as a reference text on the 7070 Sort 90 program. As such, it contains the complete documentation of the system replacing the previously published Sort 90 information contained in the bulletins, "IBM 7070 Sort 90 and Merge 91 Specifications," form J28-6040, and "IBM 7070 Sort 90 and Merge 91: Timing Estimates, Modifications," form J28-6069, and "IBM 7070 Sort 90: Operation," form J28-6096.

The information in this manual is for the most part self-explanatory. However, portions of the text presume a knowledge of the IBM 7070 Input/Output Control System (see bulletin, form J28-6033-1). Also, a knowledge of the operating principles of magnetic tape applications is presupposed. A glossary of tape terminology appears as an appendix to this manual.

Sort 90 will operate on both the IBM 7070 and IBM 7074. The use of the term "7070" in this manual implies the combined 7070/7074 system unless otherwise stated.

Program decks for 7070 Sort 90 are available in 7070 Condensed (five-instruction-per card) Load Card form; requests for decks should be addressed to:

7070 Program Librarian
IBM Corporation
590 Madison Avenue
New York 22, New York

PROGRAM DESCRIPTION

From the standpoint of execution, each of the Sort 90 phases is an independent program. Phase I is loaded and proceeds to execute itself, followed by Phase II, then Phase III. Certain information is left in 7070 storage between phases so that each phase can be properly initialized. The Sort 90 control cards prepared by the user are read and analyzed by Phase I.

Each phase contains an "assignment routine" which, based upon its interrogations of control card information and information passed on from previous phases, makes all necessary modifications and adjustments to the main portion of the phase, or "running program," which actually carries on the sorting procedure.

PHASE I

Phase I reads the records of the file to be sorted and writes them, in preparation for Phase II, as a series of long sorted sequences of records.

Phase I does not employ the serial, record-by-record processing technique used by Phases II and III and typical of most data processing applications. Instead, a group of records is read into 7070 storage; Phase I uses an internal sorting procedure to form the records into a sorted sequence. Once this has been accomplished, it writes the sequence of records. Overlap of input/output and processing occurs; while one set of records is being internally sorted, the previously sorted sequence of records is being written, and the records which are to compose the next sequence are being read. The number of records which can be formed into a sorted sequence is therefore that number which can be grouped together in 7070 storage while still overlapping input/output operations; it depends upon the length of the records and the amount of 7070 storage available to Phase I.

The sorted sequence produced as output during Phase I normally includes the records of several successive input tape blocks, and is written on tape in the form of one or more tape blocks. The blocking factor or block size used for Phase I output (and also Phase II input and output and Phase III input) is maximized according to 7070 storage availability in Phases I, II, and III, without regard for the number of records composing a Phase I output sequence; short length blocks are written freely by Phase I.

Phase I output tapes must be two, three, four, or five in number, according to the order of merge desired in Phase II. A Phase I output sequence will be written on the output tape used for the previous sequence if, by doing so, a longer combined sequence is formed. Otherwise, output tapes will be used in rotating order in an attempt to write approximately the same number of sequences on each tape.

An important feature of the Phase I internal sorting procedure is that the time required will be sharply reduced by the amount of "natural" sequencing present among the input records. This fact, combined with the treatment of output tapes described in the preceding paragraph, make the sorting of an already or partially sorted file a rapid operation.

PHASE II

Phase II, by means of repeated passes of two- , three- , four- , or five-way merging operation, makes progressive reductions in the number of sequences or records present in the file being sorted, until the number of sequences is equal to or less than the order of merge, at which point Phase III can take over. The total number of tape units it requires is equal to twice the order of merge. The input tapes to the first Phase II pass are the Phase I output tapes; the input tapes to further Phase II passes are the output tapes of the previous pass. Thus, at any one moment, half of all the tape units are used for reading, and the other half for writing.

During Phase II, the records of the first sorted sequences on the input tapes are read, merged, and written as a long combined sequence on one of the output tapes. Then, the records of the second sorted sequences on all the input tapes are read, merged, and written as a long combined sequence on another of the output tapes. The process continues until all input tapes have been exhausted, at which point the pass is concluded, and all the tapes rewound. Since the input tapes for the next pass are the output tapes of the current pass, the next pass will merge the longer combined sequences from the current pass into even longer combined sequences.

In order to know when the current sequence from each of the input tapes has ended, Phase II compares the first record of each new input tape block with the last record of the previously read tape block for each input tape; an out-of-sequence condition means that the just-read block begins a new sequence. To minimize the number of the once-per-block comparisons, and to optimize the speed of tape reading and writing, the sort blocking factor or block size determined by Phase I is made as large as possible.

PHASE III

Phase III performs the concluding merge pass, and writes the sort output file. In doing this it uses the blocking factor or block size specified by the user rather than that used during Phase II, and writes tape labels as directed by the user. It does not use the input sequence break comparison routine used in Phase II, since no input reels will contain sequence breaks; however, it sequence checks each output record with the previous output record to ensure that a truly sorted file is being produced.

Because the records are written by Phase III in sorted order, it is frequently desired to supplement Phase III with special processing routines. Accordingly, the number and range of modifications permitted in Phase III (see page 66) are particularly large.

EXAMPLE

An example of the combined operation of the three Sort 90 phases is given in Figure 1. The example assumes 12,800 input records, a Phase I output sorted sequence length of 50 records, and a four-way merge in Phases II and III. Phase I produces 256 sequences of records. Each Phase II pass reduces the number of sequences by a factor of 4 and increases the length of each sequence by a factor of 4. After three Phase II passes, Phase III can take over and complete the merging operation.

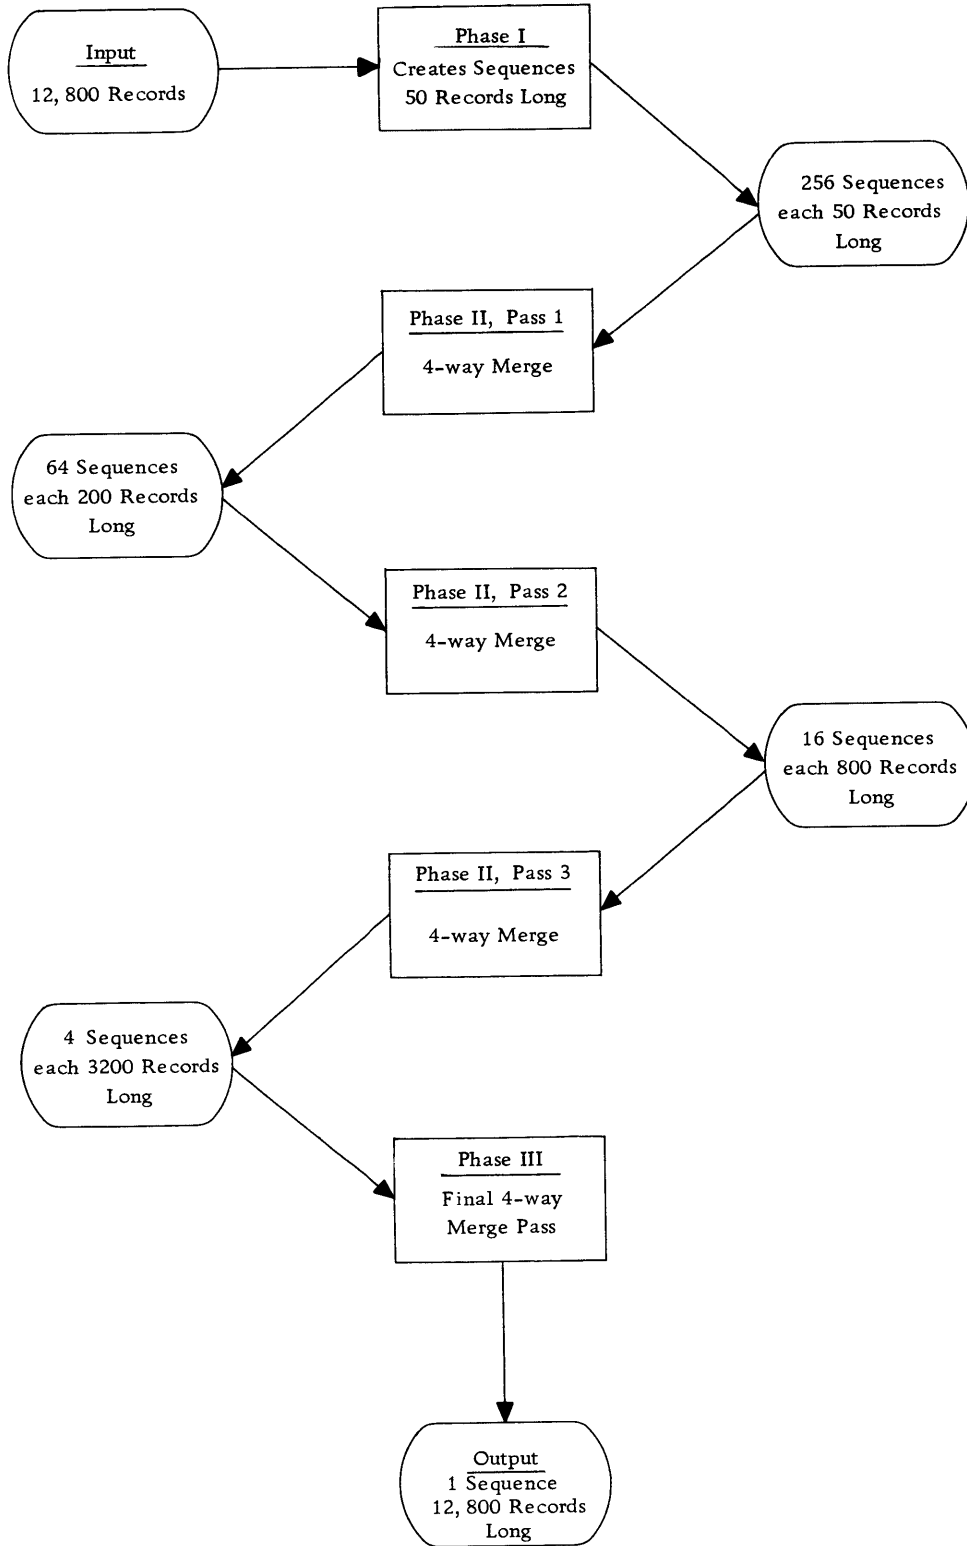


Figure 1

The key elements in determining the number of Phase II passes which will be required for a given Sort 90 application are: the number of records in the input file, the length of a Phase I output sorted sequence, and the order of merge used in Phases II and III. Expressed arithmetically, the number of Phase II passes required is:

$$\log_M \left(\frac{N}{G} \right) - 1$$

where

N is the number of records in the input file,
G is the number of records in a Phase I output sorted sequence,
M is the order of merge employed.

If the result has a fractional part, it must be increased to the next highest whole number, as a fractional part of a Phase II pass cannot be performed.

A function produced by varying one, two, or all of these factors is not smooth, but has abrupt discontinuities. If in the example shown in Figure 1, Phase I produced sorted output sequences 199 records long, three Phase II passes would still be required. On the other hand, a G of 200 would result in only two Phase II passes, and a G of 49 in four Phase II passes.

If G remained at 50, an input file of 12,801 records would require four Phase II passes; but not until the input file contained 3200 records would only two Phase II passes be required.

The longer the sorted sequence produced by Phase I, the longer the processing time required per record. It is generally advisable, therefore, to cause G to be less than its maximum possible value, as long as by reducing it the number of Phase II passes required does not increase. Ways of controlling the length of a Phase I sorted output sequence are described under Phase I Area Allocation.

The user may wish to make some trial runs, or to spend some time over the timing procedures, in order to determine the optional use of Sort 90 for a given sorting application.

TAPE ALLOCATION AND HANDLING

As previously mentioned, two sets of work tape units are required during all the passes of Phase II of Sort 90. In any one pass, input records are contained on one set of work tape units, and output records are written on the other. At the beginning of the next Phase II pass, the roles of the two sets of work tape units are reversed.

The two sets of work tape units will be called the "A-units" and the "B-units." When Phase II begins, its input (i. e. , the output from Phase I) will be contained on the B-units; hence, the output from the first pass of Phase II will be written on the A-units.

The number of A-units and the number of B-units should be identical. The combined total number of units required is equal to twice the order of merge for a given application of Sort 90. The A-units and B-units will normally be mounted on separate tape channels, unless only one is available. A picture of A-units and B-units for a two-way Phase II merging operation is illustrated in Figure 2(a).


INPUT TAPE UNITS

At least one tape unit is required for reading the sort input file in Phase I. If the input file is contained on more than one reel of tape, the use of two or more tape units for Phase I input will be found convenient so that the operator may "flip-flop" input reels in order to avoid interrupting Phase I for reel-changing. Tape units used for mounting the sort input will be called "I-units."


Reels of the input file must be mounted in rotation on the specified I-units, which may be as many as five in number. Thus, if I-units were designated as tape units 4, 0, and 8, on channel 2, reels of the input file should be mounted as follows:

<u>Unit 24</u>	<u>Unit 20</u>	<u>Unit 28</u>
Reel 1	Reel 2	Reel 3
Reel 4	Reel 5	Reel 6
Reel 7	Reel 8	Etc.


It should be understood that when Phase I begins reading reel 2, the operator should dismount from Unit 24, reel 1 and mount reel 4; when Phase I goes on to reel 3, the operator should dismount from Unit 20, reel 2 and mount reel 5; and so forth. Reel 1 must have been mounted before Phase I begins its execution, but reels 2 and 3 may be mounted while Phase I is reading reel 1. Exhausted input reels will be rewound-unloaded.




(a)



(b)



(c)



(d)

Figure 2

INPUT TAPE UNITS VS. WORK TAPE UNITS

It is not necessary that I-units be unique tape units used only for the reading of the sort input file. I-units and B-units must be wholly separate, however, since B-units are used for Phase I output. Three different relations are possible between I-units and A-units; they are described below in order of increasing operator convenience and increasing number of tape units.

The I-units and A-units may be identical as shown in Figure 2 (b). When Phase I has read the last input reel to be mounted on a given I-unit, the operator must mount a work tape in preparation for Phase II. When the last input reel of the input file has been read, the operator must pause between Phase I and Phase II to mount a work tape on the last tape unit used.

The I-units and A-units may be partially identical and partially separate as shown in Figure 2 (c). When Phase I has read the last input reel to be mounted on any shared tape unit, the operator must mount a work tape in preparation for Phase II. However, input reel mounting can be planned so that the last input reel is read from a separate I-unit, removing the need for a reel-changing between Phase I and Phase II. Mounting of work tapes on separate A-units in preparation for Phase II may be done at any time during Phase I.

The I-units and A-units may be wholly separate, as shown in Figure 2 (d). In this case, no work tapes need ever be mounted on I-units; mounting of work tapes on A-units in preparation for Phase II may be done at any time during the execution of Phase I.

OUTPUT TAPE UNITS

Either the A-units or the B-units will contain the sort output file written by Phase III. If the number of Phase II passes is even, or if no Phase II passes are preformed, A-units will be used for writing Phase III output; if the number of Phase II passes is odd, B-units will be used for writing Phase III output.

The low-order output tape unit, i. e. , the last-punched in the relevant field of Control Card 1 (see page 27) will be used during Phase III to hold a checkpoint record. The remaining A- or B-units will be used for actual writing of the output file. Thus, one tape unit is available for writing output when performing a two-way merge; two, with a three-way merge; three, with a four-way merge; and four, with a five-way merge.

Output tape units will be used in a rotating fashion similar to that described for input tape units (see page 8). A completed output reel will be rewound-unloaded, and typewriter messages will direct the operator to dismount the reel.

EFFECTS OF NUMBER OF TAPE CHANNELS

A maximum of two tape channels may be used during the sort program. If only one tape channel is available, I-units, A-units, and B-units must, of course, all be mounted on that one channel; the relation between I-units

and A-units may be any one of the three described above, and reel-changing considerations are not affected. Tape reading and tape writing will be completely unoverlapped in all passes of Sort 90.

If two tape channels are available, the I-units and A-units should be on one channel, and B-units on the other. I-units and B-units may not be on the same channel. It is possible, however, for A-units and B-units to be on one channel, and I-units on the other; in this case the advantage of overlapping tape reading and tape writing in Phases II and III is lost, and Sort 90 will come to a precautionary halt (Halt 1003, see page 40) should such a configuration be elected.

The actual number of the channel to which any set of tape units is affixed is irrelevant to the logic of Sort 90 and may be channel 1, or channel 2, at the user's option.

MISCELLANEOUS TAPE UNITS

Dump Tape

If a spare tape unit is allocated for the purpose of "dumping" unreadable records, the tape unit used may be any tape unit on channel 1 or 2.

Program Tape

If the Sort 90 program is loaded from tape, the tape unit used may be any tape unit on any tape channel.

INPUT FILE CONSIDERATIONS

RECORD TYPES

Input files to be sorted by Sort 90 may be composed of Forms 1, 2, or 3 records, these record forms having the meaning given to them in the 7070 Input/Output Control System. All records of any one input file must be of the same record form. The information about the record which must be supplied as control information to Sort 90 relates to data record length (or maximum length) in words, and blocking factor or block size. In addition, the user is asked to specify for Form 3 files the length, in words, of the shortest record in the file. It is very important that this figure be accurate; too large a figure will cause an error stop when a shorter record is detected in Phase I, too small a figure will lead to inefficiencies in the operation of Phase I.

Sort 90 is written to accept short-length blocks in the input file, regardless of record form, and will correctly process the short-length blocks. Detection of a segment mark in the input file will cause the program to come to a halt (Halt 1126); if the Start key is depressed, the program will disregard the segment mark and continue reading the file. Both long-length blocks and short-character blocks will cause similar halts (Halts 1124 and 1127), after which depression of the Start key will cause the record involved, or as many words from it as are entered in 7070 storage, to be typed; the program will then disregard the block and continue reading the file.

RECORD AND BLOCK SIZES

Because of the number of affecting factors, it is difficult to make a precise determination of the maximum data record lengths, tape block sizes, and blocking factors acceptable in a given application of Sort 90. The maximum permissible Form 1 or Form 2 input blocking factor is the Phase I output sequence length taken from either Table I or Table II in Appendix II. The maximum permissible Form 1 or Form 2 output blocking factor is approximately equal to

$$\frac{H-2000-150M-2ML}{2L+2}$$

where

H is the Phase III storage limit (as punched in columns 52-55 of Control Card 1),

M is the order of merge being employed (2, 3, 4, or 5), and

L is the (maximum) number of words in a data record.

The maximum permissible Form 1 or Form 2 data record length, and the maximum permissible Form 3 tape block size, are approximately equal to

$$\frac{J-2350-150M}{2M+2}$$

where

J is the Phase II storage limit (punched in columns 48 through 51 of Control Card 1) and
M is as defined above.

The figures resulting from these expressions may not exceed 999.

NUMBER OF RECORDS ALLOWABLE

The maximum number of records which may be sorted in any one application of Sort 90 depends on a number of variables, the most important being the order of merge. The maximum allowable number of records is that number which, when written (1) at high density, (2) with zero suppression, and (3) with the blocking factor used by the sort in Phase II, may be contained on a number of reels of tape equal to one less than the order of merge being employed. Thus,

- A two-way sort may process one full reel.
- A three-way sort may process two full reels.
- A four-way sort may process three full reels.
- A five-way sort may process four full reels.

Expressed in a formula, the maximum number of records which may be sorted in one application of Sort 90 is:

$$\frac{15,345,600 B (M-1)}{BC+417}$$

where

B is the sort blocking factor (for Form 1 and 2 records, this figure is typed out by Phase I; for Form 3 records, the figure typed by Phase I should be divided by the average data record length),
M is the order of merge being employed (2, 3, 4 or 5) and
C is the number of tape characters produced by writing a data record (of average length) with zero-suppression.

In this formula, a reel of tape is assumed to contain 2300 feet of usable tape.

An attempt to sort more than the maximum allowable number of records is very dangerous; a situation may result in the late passes of Phase II such that the sort program cannot possibly continue on to completion. The file may be broken into sections, each section sorted independently, and the resultant sorted sections merged together; the time expended on the initial uncompleted sort is however completely lost. A complete explanation of this "non-ending" sort situation is given in Appendix III.

NUMBER OF REELS ALLOWABLE

A Sort 90 input file may be contained on a large number of physical tape reels. Although the maximum permissible number of records should not be exceeded, the reels containing the file need not each be completely full.

Detection of an EOF trailer label in the last reel of the sort input file will signal to Sort 90 the termination of input. In addition, two columns of the Sort 90 first control card are provided for punching the number of reels composing the input file. A programmed halt will occur if this number of reels has not been read when the EOF trailer label is detected, or if this number of reels have been read and the EOF trailer not yet detected. In both cases the operator may elect either to continue reading input reels or to go to Phase II.

Use of the reel-counting feature is optional. When it is used, the maximum number of reels which may compose the input file is 99. When it is not used, with the trailer label EOF indication signaling the last input reel, the number of reels which may compose the sort input file is unlimited. If tape labels are not used in the sort input file, the reel-counting feature must be used and is the sole determinant of the number of reels composing the sort input file.

CONTROL DATA SPECIFICATIONS

The selection of the data upon which one record is to be compared with another during the operation of Sort 90 is completely controllable by the user of the program. Sort 90 Control Card 2 is punched with starting digit positions and lengths of all desired control data fields, which collectively make up the control data for a given application of Sort 90. The ranking of the control data fields (major through minor) is accomplished by the order in which their descriptions are punched in the control card.

CONTROL DATA FIELDS AND SEGMENTS

Up to ten control data fields will be accepted by Sort 90 for use in controlling the order of records. Control data fields are described by their digit positions in a record of the file being sorted, the record being pictured as it lies in 7070 memory after having been read from tape or as laid out by a 7070 Autocoder "DA" entry. A control data field may begin in any digit of the record and extend over any number of succeeding digits, bridging 7070 word boundaries as necessary. Figure 3 illustrates a possible pattern of control data fields.

Line	Label	Operation	OPERAND
3	5/6	15/16	20/21 25 30 35 40
0 1		DA	1
0 2	CNTRLFLD1		4 0 . . 4 9
0 3	CNTRLFLD2		0 5 . . 2 3
0 4	CNTRLFLD3		0 3
0 5	CNTRLFLD4		2 4 . . 2 6
0 6	CNTRLFLD5		2 8 . . 3 9
0 7	CNTRLFLD6		0 1 . . 0 2
0 8			

Figure 3

If two or more control data fields occupy contiguous left-to-right positions in the record and are in descending control importance (i. e. , major control field at the far left and minor control field at the far right), they should be considered and specified to Sort 90 as a single control data field.

Sort 90 will not alter the control data of any record. That is, it will not rearrange the control data of records being merged so as to "pack" control data together. In setting up its compare routines, however, it breaks down the control data field descriptions given in the control card into control data segment descriptions. A control data segment is that portion of a control data field which lies within boundaries of a 7070 storage word. The number of control data segments comprising a control data field is, therefore, equal to the number of 7070 words over which the field extends. In Figure 3, CNTRLFLD2 is made up of three control data segments, whereas CNTRLFLD4 is made up of a single control data segment.

The number of control data segments contained in the control data fields specified for a given application of Sort 90 may not exceed sixteen in number.

The two restrictions that (1) there be no more than ten control data fields and that (2) these control data fields be composed of not more than sixteen control data segments constitute the only restrictions upon control data for Sort 90. It can be seen that the maximum permissible number of digits of control data is 160, a figure which provides for many different configurations of control data. The control data of the record described in Figure 3 consists of six fields, nine segments, and 47 digits.

The Phase I assignment routine will come to a programmed halt should any of the control fields described on the control card fall outside of the indicated record length. If two or more control fields overlap one another in the record, messages to this effect will be typed, followed by a programmed halt; the user may continue the program if the overlapping is desired.

It should be noted that sorting upon one set of control data destroys any other sequencing in the file being sorted. If a file were ordered upon a field and it were desired to put it into major or minor order upon a second field, it would not be sufficient to sort on the second field. The control data must consist of both fields, specifying the desired major-minor relation. Furthermore, Sort 90 will not necessarily maintain the relative input order of a group of records with identical control data; the records composing such a group will be consecutive in the output file, but in whatever order within the group the program finds it most convenient to place them. In these two instances, the procedures followed by Sort 90 differ from the procedures possible with most mechanical sorting devices.

OPTIMAL CONTROL DATA ARRANGEMENT

The exact configuration of control data fields, especially of the more major ones, can influence the Sort 90 process time. The sort comparing networks order two records by comparing their control data, one control segment at a time, beginning with the most major segment and continuing until a difference has been found, or through the most minor control data segment if the records have equal control data. The number of separate comparisons may be minimized if the number of control data segments is minimized, and this in turn requires that each control data segment contain as many digits as possible up to ten (a full word). This factor should be considered when designing the format of records which will be sorted frequently. It is more important for major control data fields than for minor control data fields, because (except for files composed of many records with equal control data) the minor control data fields will seldom enter into comparison.

The procedures for calculating Sort 90 running times (see page 76) require that the user make an estimate of the number of control data segments, beyond the first, which will be involved in the comparison of two "average" records of the file being sorted. In the example given in Figure 3, if it were thought that comparisons of two records would usually involve CNTRLFLD4 and almost never CNTRLFLD5, then the number of control data segments, beyond the first, involved in a typical comparison would be five in number.

COLLATING SEQUENCES

Sort 90 permits the specification of any one of four different collating sequences to govern the procedure by which it orders records: ascending and descending algebraic, and ascending and descending absolute. The normal procedure for 7070 compare operations is the basis of all the collating sequences. In an ascending algebraic collating sequence, a record whose control is lower than that of a second record will precede the second record in the sort output file; with a descending algebraic collating sequence, just the reverse. In both cases, however, the signs of the words composing the control data enter into the comparisons. For example, +0 is greater than -0. The ascending and descending absolute collating sequences are respectively identical, but the signs of the words composing the control data are ignored.

Figure 4 illustrates six records after being sorted into an input file. The example shows these six records as ordered according to the four different collating sequences. It assumes that six one-word records have been sorted, with the entire record consisting of control data as follows: % TAX (double-digit @3600836187), -5555555555, +8888888888, DEDCT (double-digit @6463646283), -1111111111, +4444444444. The first output record of each collating sequence is shown at the top of each list.

Ascending Algebraic	Descending Algebraic	Ascending Absolute	Descending Absolute
% TAX	+8888888888	-1111111111	+8888888888
DEDCT	+4444444444	% TAX	DEDCT
-5555555555	-1111111111	+4444444444	-5555555555
-1111111111	-5555555555	-5555555555	+4444444444
+4444444444	DEDCT	DEDCT	% TAX
+8888888888	% TAX	+8888888888	-1111111111

Figure 4

In this example, the control data is alphameric in some records and numerical in others. Such a situation does not represent typical practice, but may arise in some instances in which the sign of a control data field is used as an indicator or switch rather than as a designator of the actual sign of the control data.

MISCELLANEOUS FEATURES AND PROVISIONS

RECORD COUNTS

Sort 90 will maintain in each pass of its operation a count of the records processed in that pass. At the end of the pass the count will be compared with that taken in the preceding pass. A discrepancy will cause a programmed halt, at which time the operator may discontinue the sort, re-execute the current pass, or disregard the discrepancy and continue the sort.

The record count accumulated in Phase I may be checked against a record count (if present) contained in the trailer tape label of each reel of the input file. A discrepancy will cause a programmed halt, at which time the operator may (1) disregard the discrepancy and continue Phase I, (2) cause the records of the current reel to be dropped from the sort and continue Phase I, or (3) cause the current reel to be reprocessed. Phase III will, if desired, insert a count of the sort output records into the trailer tape label of each output reel.

Record counts in Phases I and III are accumulated on a reel-by-reel basis, to permit trailer tape label checking and preparing. The reel-by-reel counts are totaled to permit the pass-by-pass record count checking. If a record count is present in a trailer tape label, it is assumed to occupy columns 11-20 of the label. When read into the machine, two words of double-digit information, with the low-order character unsigned, are produced.

The user need not consider the effects upon record counts of modifications which cause the number of records being processed to vary from pass to pass (deleting, inserting, and summarizing). Sort 90 will automatically reconcile record counts when these modifications are present.

The record count procedure to be performed by Sort 90 is specified by the user through control card punching (see page 29).

HASH TOTALS

An optional feature of Sort 90 is the accumulation of a "hash total," which is a control total maintained upon a selectable field of each record of the file being sorted, often a field other than an amount field, such as part number.

The hash total, if elected, is taken in a fashion similar to the taking of record counts (see Record Counts above). The checking features offered are completely identical to those offered with record counts, with the exception that a hash total pass-by-pass check is optional rather than standard.

If a hash total is present in a trailer tape label, it is assumed to occupy columns 21-30 of the label. When read into the machine, two words of double-digit information, with the low-order character unsigned, are produced.

A field upon which a hash total is to be taken must be contained within one word of the data record. Its sign is assumed plus, although the actual sign of the word containing the field may be plus, minus, or alpha. The hash total itself is maintained as a ten-digit positive numerical field, carries beyond the tenth digit being disregarded.

In each pass of the sort in which a hash total is taken, the basic process time of the pass is increased by 0.21 seconds per 1,000 records. In such passes, Sort 90 will operate in the "sense" mode as regards field overflow; it will operate in the "halt" mode in passes in which a hash total is not taken.

The hash total action to be performed by Sort 90 is specified through control card punching (see page 29).

PHASE I AREA ALLOCATION

In its reading, processing, and writing of records, Phase I is capable of using either a three-area or two-area system. The number of records which Phase I will produce as a sorted sequence is that number which will fit in one area. Under the three-area system, tape reading, tape writing and processing are all overlapped. At any given instant records are being read into one area, processed in the second area, and written from the third. During the next cycle, records are read into the third area, processed in the first, and written from the second area. Both tape reading and tape writing are thus being performed concurrently with the processing of records and with each other.

Under the two-area system, tape reading and writing overlap processing; however, tape reading and writing are not overlapped. Records are first written from, and then read into, one area while processing is being performed in the second area. When the next cycle begins, the processed records are written from the second area, which is then filled with new records, while processing is being performed in the first area.

The two-area system is offered because the Phase I process time will frequently exceed the sum of the tape-read and tape-write time. If so, the overlapping of tape reading and tape writing obtained with a three-area system will be valueless, whereas the two-area system will permit a Phase I sorted sequence to contain approximately one and one-half times the number of records possible with a three-area system.

Through control card punching (see page 30), Sort 90 may be directed to use a three-area system, a two-area system, or itself compute and use the better of the two systems.

The third option makes use of an extensive computation and selection routine with the Phase I assignment routine. Successful operation requires the knowledge of certain parameters which are not contained in the Sort 90 control cards. Modification I-2 (see page 56) should be consulted in this regard before use of the "compute" option in Phase I.

If Control Card 1 indicates that I-units and B-units share a single tape

channel, the area allocation option indicated by the user will be disregarded and a two-area system selected automatically.

CHECKPOINT AND RESTART

Two different types of checkpointing procedure are used in Sort 90: a specialized one in Phase I, and the IOCS procedure in Phases II and III.

In Phase I, segment marks are written on all output tapes whenever the processing of a new input reel is about to begin. If after the processing of an input reel, discrepancies of any kind occur, the operator may direct the program to either reprocess the input reel or to drop all records of the input reel from the sort. In the former the input reel is rewound, and the output reels are repositioned with Tape Segment Backward Space commands. In the latter Phase I proceeds on to the next reel of the input file, after repositioning the output reels with Tape Segment Backward Space commands. The segment marks thus left on the Phase I output reels will be ignored when read by Phase II.

Interruption and restart in Phases II and III is accomplished by the writing of the IOCS restart program and checkpoint records on the low-order (last-punched in Control Card 1) tape units of both work tape channels. Phase I writes the restart program on the low-order Channel B tape unit, and Phase II writes it on the low-order Channel A tape unit. At the beginning of each Phase II pass and of Phase III, the restart program contained on the low-order output tape is spaced over and a checkpoint record is written after it, and normal Sort 90 processing commences. By virtue of this procedure, any pass may be interrupted at any point, and the sort later restarted from the beginning of the pass during which it was interrupted. Details of the checkpoint and restart procedure are given under Operation (see page 24).

It is important to note that the checkpoints in Phases II and III permit the user completely to discontinue the sort and to dismount the input and output reels, whereas the restart procedure in Phase I will not. The earliest point at which the Sort 90 program may be interrupted without having to be restarted from the beginning is a few seconds after the start of the first Phase II pass.

UNREADABLE RECORDS PROCEDURE

The unreadable records procedures offered by Sort 90 are identical to those offered by the 7070 Input/Output Control System. Through control card punching, the user may cause a programmed halt, typing of the block involved followed by a programmed halt, dumping of the block involved on a separate "dump" tape, or continuous retry. If a "dump" tape is furnished, it will be assumed to have no header label, and will not be affected by the checkpoint and restart provisions of Sort 90; therefore, there is the possibility that if a pass during which a read error occurs is interrupted and later restarted, the erroneous block will appear twice on the dump tape.

If an attempt is made to replace through console actions a word containing invalid digits, and the word should be a control data field of a record,

storing the incorrect contents will not be immediately discovered by Sort 90 and may result in a sequence break in Phase III. Improper correction of other fields may result in other malfunctions of Sort 90. This option should therefore be used with discretion.

SORTING TAPE FILES PREPARED BY

OTHER IBM DATA PROCESSING SYSTEMS

In most cases, the generality of Sort 90 will permit its use in sorting tape files prepared by IBM Data Processing Systems other than the 7070. These include the 650, the 705 (all models) and 7080, the 704, 709, and 7090 (using the BCD mode), and the 1401.

The 1401 is capable of writing the mode change character (CB8421) on tape, and can thus prepare files which the 7070 can read in alphamerical form (including zero-suppression in numerical words). The other systems, however, cannot normally write the mode change character, and must prepare records which to the 7070 are completely alphabetic and a multiple of five characters in length. Any 650 word is written as five or ten characters, and the restriction is thus met for the 650. Also, records prepared by the 705/7080 will be multiples of five characters in length if they have been designed for 705/7080 internal high-speed transmission. Since 704/709/7090 records must be a multiple of six BCD characters in length, only those 704/709/7090 records which are thirty, sixty, ninety, etc., 7070 characters in length may be processed by the 7070.

In all cases, the record and block structure of the input tape file must agree with the 7070 Input/Output Control System specifications for record Forms 1, 2, or 3.

The collating sequence of the 7070 is identical to those of the 705/7080, 704/709/7090, and 1401, and applies also to alphabetic information in 650 records. The role of the sign in 650 numerical words, however, is lost when the 650 word is read into the 7070 as two alphabetic words, and will not sort 650 records into the same order as would a 650 sort if some of the control data is numerical. The difficulty may be overcome in this case by adding editing instructions to Phase I of Sort 90 (see Modification I-6) to convert 650 numerical words to 7070 numerical form, and complementary editing instructions to Phase III to reconvert the words to 7070 alphabetic form so that an output record will be identical to the original input record.

The remaining area of possible difficulty is tape labels. If files prepared by other computers to be sorted by the 7070 have tape labels, it is suggested that the user consult the section Tape Labels and Modifications I-8 and III-4, as well as the 7070 Input/Output Control System specifications for tape labels.

TAPE LABELS

Sort 90 includes full provisions for processing of tape labels of the standard format used by the 7070 Input/Output Control System. Through control card punching the input file, work tapes, and output file, respectively, are

indicated as labeled or unlabeled. Control cards are furnished for entering information to be used in performing input file tape label checking and in establishing the contents of output file tape labels. Modification I-8, for the input file, and Modification III-4, for the output file, should be consulted if it is desired to alter the options for tape label processing which are assumed by Sort 90. These options call for full checking of tape labels, and assume that if output file tape labels are specified, the work tapes, some of which will be used for writing the output file, are labeled.

If work tape labels are specified, new labels will be written by Sort 90 after the existing labels have been checked to ensure that the work tapes are available for writing. The contents of a work tape label will be:

Tape Serial Number:	The Tape Serial Number from the existing tape label.
File Serial Number:	The Tape Serial Number from the existing tape label on the low-order (last-punched) Channel B tape unit.
Reel Sequence Number:	The digits 0CU, where C is the channel number and U the unit number of the tape unit on which the work tape is mounted.
File Identification:	SRT90WKT Pb
Creation Data:	Today's date as stored in 7070 location 0109.
Retention Cycle:	000

OPERATION

The operating procedure for Sort 90 is straightforward, involving primarily loading the program, mounting tape reels promptly and on appropriate tape units, and (when necessary) interrupting the program and later restarting it. Operating instructions are given below for initiating Sort 90, for actions to be performed during execution of the program, and for interrupting and later restarting the program. A final section discusses the settings of console switches and dials.

The Sort 90 program deck is furnished in 7070 Condensed Load Card form, and may be loaded from a card reader or a tape unit by the 7070 Condensed Card Load Program (not furnished with the Sort 90 program deck). Other loading procedures may be used as desired; consult Modifications I-2, I-10, and II-2 as to entry points for loading procedures. Modification III-6 should be examined in connection with end-of-job procedures, especially if SPOOL applications are to be performed concurrently with Sort 90.

INITIATION OF SORT 90

1. Ready (at least) the first reel of the sort input file on the first tape unit of those allocated to the input file.
2. Ready work tapes on (at least) all tape units allocated as Channel B tape units.
3. Ensure that the Sort 90 Control Cards appropriate to this application have been furnished and that they are correctly positioned in whatever input unit (see Control Card Reading) is to be used for reading them.
4. Ready the Sort 90 program in the input device to be used for loading it, and initiate the loading.

ACTIONS DURING SORT 90

Under normal operation of Sort 90, actions to be performed during its execution concern the handling of input, work, and output tape reels:

1. During Phase I, mount successive reels of the input file on the tape units allocated to the input file, always attempting to have the next reel of the input file ready for processing before it is called upon by Phase I.
2. Also during Phase I, mount work tapes on Channel A work tape units which are not also input tape units. When the processing of the last input reel mounted on an input tape unit which is also a Channel A tape unit is complete, ready a work tape on the unit.
3. During Phase III, remove and visually label successive reels of the output file after their processing is complete, unless the output file is to be used as input by the program which follows Sort 90.
4. During all phases of Sort 90, if desired and convenient, prepare control cards, card decks, and tape units and files to be used by the program which will follow Sort 90, so that the set-up time after Sort 90 concludes will be as short as possible. Be careful, however, in rotating the address selection dials on tape units not used by Sort 90, lest they be moved even momentarily through positions used by sort tapes and thus interfere with the correct operation of Sort 90.

INTERRUPTION AND RESTART

At the beginning of each Phase II pass, and at the beginning of Phase III, a message will be typed of the form

CHPT +CUxxxxxNNN

The message indicates that a checkpoint record is about to be written; CU is the channel and unit number of the tape unit upon which the checkpoint record will be written, and NNN is a serial number identifying this particular checkpoint record. An interruption and later restart always involves the last-typed of these CHPT messages.

WHEN TO INTERRUPT

Sort 90 may not be interrupted until the first CHPT message has been typed at the beginning of Phase II. Subject to this restriction, it may be interrupted at any of the programmed halts which may occur during the execution of Phases II and III, or may be interrupted merely by depressing the Stop key on the 7070 console. In either case, if SPOOL programs are being executed concurrently with Sort 90, the appropriate procedure for continuing their execution should then be followed.

The user should not interrupt Sort 90 during and for several seconds after the typing of a CHPT message. A checkpoint record is being written during this period, and interference with its writing will make restarting impossible. It is also inadvisable, although permissible, to interrupt Sort 90 when the end of a tape pass is imminent, or when some or all of the work tapes are rewinding. A new checkpoint record will soon be written, and Sort 90 processing after the restart has been performed will be much more efficient if that checkpoint record, rather than the previously written one, is employed.

ACTIONS AFTER INTERRUPTING

After interrupting Sort 90, all work tapes should be rewound manually. No further action is required if the restart is to be effected immediately; but if the restart is to be delayed for a period of time these additional steps should be performed:

1. Remove and save the console typewriter log.
2. Remove and visually label with the address of the tape unit from which each was removed the following tape reels:
 - a. The work tape on the tape unit designated by the last-typed CHPT message.
 - b. All work tapes on the other tape channel: i. e. , all work tapes on Channel A, if the tape unit designated by the last-typed CHPT message is on Channel B; or all work tapes on Channel B, if the latter tape unit is on Channel A.

RESTART PROCEDURE

If the restart is to be effected immediately after interruption, perform only

step 4 below. Otherwise perform these actions:

1. Mount on the tape units from which they were removed all work tapes removed after interruption of the sort.
2. Mount work tapes on all other work tape units.
3. If the interruption was made during Phase III, go to step 4. If it was made during Phase II, ensure that the Sort 90 program is correctly positioned in the input device being used to load it. The Sort 90 program may be positioned immediately before either Phase I, Phase II, or Phase III; it will itself locate and load Phase III when Phase II is completed after the restart.
4. Follow the 7070 Input/Output Control System restart procedure. The Restart Initiator routine used in this procedure may be assembled in any 7070 locations below location 0325 or above location 4200. A description of the 7070 Input/Output Control System restart procedure is not given here; however, the key steps are as follows:
 - a. Load the Restart Initiator.
 - b. At Halt 0050, store into 7070 location 0050 the numerical word from the last-typed CHPT message, and then depress the Start key.
 - c. At the message RESTORE CONSOLE... and associated halt, ensure that console dials are correctly positioned and relevant input/output units are ready, and then depress the Start key. If the load program was altered in order to load the Restart Initiator, restore it to its previous condition.

NOTES AND CAUTIONS

The following points should be noted in connection with interruption and restart of Sort 90:

1. The manual rewinding of work tapes mentioned in Actions after Interrupting is not necessary for a successful restart, but will save a considerable amount of time after the restart has been performed and should be observed as a desirable operating practice.
2. If an interruption is made in Phase III, and tape labels are specified for the sort output file, the header tape labels on all tape reels upon which output has been written up to the point of interruption will contain today's date as Creation Date, and the number punched in Control Card 4 as Retention Cycle. If these reels are later used for the writing of some output file (by Sort 90, after restart or by any other program) it is likely that the Input/Output Control System will consider them unavailable for output purposes. So that this action may be correctly ignored by the operator, it is recommended that these tape reels be visually labeled as being available for output purposes at the time of interruption of Sort 90.

CONSOLE SWITCHES AND DIALS

During operation of Sort 90, switches and dials on the 7070 Console should normally be set as follows:

Status Key - RUN

Unit Record Priority Dials - OFF and N, except as other settings are required for SPOOL operation.

Alteration Switches

Switch 1 - This switch should be set ON or OFF as appropriate after Sort 90 programmed halts.

Switch 2 - Irrelevant; not used by Sort 90.

Switch 3 - Irrelevant; not used by Sort 90.

Switch 4 - OFF

Accumulator Overflow Switches - OFF, unless required ON by modifications to Sort 90.

Exponent Overflow Switch - OFF, unless required ON by modifications to Sort 90.

CONTROL CARD PREPARATION

Control information must be provided to Sort 90 so that the program may initialize itself for the specific application desired. This information is supplied to the program by means of from two to four control cards. A minimum of two control cards is required for all sort operations. In addition, an Input File Label Control Card and/or an Output File Label Control Card (Control Cards 3 and 4 respectively) must be furnished if the reels making up the sort input and/or output files contain or are to receive labels.

CONTROL CARD READING

Sort 90 control card information is read into 7070 storage by Phase I of the program. The control cards may be read through a 7500 Card Reader mounted to any synchronizer, through a 7501 Console Card Reader or from a tape unit. As furnished, Sort 90 expects the control cards to be read through a 7500 Card Reader connected to synchronizer 1. To alter the read-in device, consult Modification I-1, discussed on page 55. The control card information may also be placed in storage prior to the beginning of the program. If so, consult Modification I-1 for altering Sort 90 so as not to read in the control cards.

CONTROL CARD CHECKING

It is important that care be taken in specifying the control card information correctly. Phase I will check the control card information extensively for consistency and reasonability, but it cannot check for absolute accuracy. Thus Phase I would come to an error stop if the control information indicated that five-word records had control data in the sixth word, or if record length were described as zero. However, Phase I could not detect a discrepancy between an indicated blocking factor of five and an actual blocking factor of six.

CONTROL CARD LISTING

Each control card must be punched in the standard load-card format; i. e. , unused columns must be filled with zeros and columns 10, 20, 30, 40, 50, 60, 70, and 80 must contain a 12 punch. If a 7500 Card Reader is used for reading the control cards, the cards should also contain a 12 punch in whatever column is used to indicate a load card.

Control Card 1

<u>Columns</u>	<u>Contents</u>
1	<u>Input Channel.</u> Number of tape channel to be used for reading the input file.
2-6	<u>Input Tape Units.</u> Numbers of the tape units on the Input Channel to be used for mounting the reels making up the input file. The tape unit numbers should be punched con-

Columns

Contents

secutively beginning with column 2; unused columns should be filled with zeros. Note that tape unit 0 may not be punched last; for example, if tape units 0, 1, and 2 are to be used, punch 01200, 02100, or 10200, but not 12000.

- 7 Input Tape Unit Type. Punch a one-digit code indicating the type of input tape units (columns 1-6) and the density with which the input file is written:
 1 - 729 II, low
 2 - 729 II, high
 3 - 729 IV, low
 4 - 729 IV, high
- 8 Channel A. Number of the tape channel to which one set of work tape units is to be attached. This number should be the same as that punched in column 1.
- 9-13 Channel A Tape Units. Numbers of the tape units on Channel A to be used for mounting work tapes. Punched in the same manner as columns 2-6.
- 14 Channel B. Number of the tape channel to which a second set of work tapes is to be attached. Unless only one tape channel is available, this number should differ from that punched in columns 1 and 8.
- 15-19 Channel B Tape Units. Numbers of the tape units on Channel B to be used for mounting work tapes. Punched in the same manner as columns 2-6.
- 20 Output Tape Unit Type. Punch a one-digit code indicating the type of work tape units (columns 8-13 and 14-19) and the density to be used in writing the sorted output file:
 1 - 729 II, low
 2 - 729 II, high
 3 - 729 IV, low
 4 - 729 IV, high
- 21 Record Form. Form of the input data records, i. e. , 1, 2, or 3.
- 22-24 Input Record Length. For record Form 1, punch the number of 7070 words composing an input data record. For record Form 2, punch the maximum number of words composing a record. For record Form 3, punch the minimum number of words composing a record.
- 25-27 Input Blocking. For record Forms 1 and 2, punch the input blocking factor. For record Form 3, punch the maximum number of 7070 words composing an input tape block.

<u>Columns</u>	<u>Contents</u>
28-30	<u>Output Blocking</u> . For record Forms 1 and 2, punch the output blocking factor. For record Form 3, punch the maximum number of 7070 words composing an output tape block.
31	<u>Zero Suppression Indicator</u> . Punch the digit 0 if the output file is to be written without zero suppression, and the digit 1 if it is to be written with zero suppression.
32-33	<u>Input Reel Count</u> . If the input tape reels are labeled, an additional check may be obtained by punching these columns with the number of reels making up the input file; if the additional check is not desired, punch zeros. If the input reels are not labeled, the number of reels making up the input file must be punched.
34	<u>Record Count Indicator</u> . Punch a one-digit code indicating the record count actions to be performed in addition to the pass-to-pass check: <ul style="list-style-type: none"> 0 - No further action 1 - Check the record count in the input file trailer labels 2 - Insert the record count in the output file trailer labels. 3 - (1) and (2)
35	<u>Hash Total Indicator</u> . Punch a one-digit code indicating the hash total actions to be performed: <ul style="list-style-type: none"> 0 - None 1 - Check the hash total in the input trailer labels 2 - Make a pass-to-pass hash total check 3 - Insert the hash total in the output file trailer labels 4 - (1) and (2) 5 - (1) and (3) 6 - (2) and (3) 7 - (1), (2), and (3)
36-40	<u>Hash Total Field</u> . If column 35 contains a zero, punch five zeros. If column 35 contains a non-zero digit, punch the position within the record of the field upon which the hash total is to be accumulated: columns 36-38 contain the word number, assuming that the first word of the record is numbered 000, and columns 39-40 contain field control.
41-42	<u>Unreadable Record Indicator</u> . Punch a two-digit code indicating the action to be taken upon detection of an unreadable tape block: <ul style="list-style-type: none"> 00 - Halt to allow manual correction 10-29 - Write the block on the tape reel mounted on this channel and tape unit

Columns

Contents

- 50 - Type the block, then halt to allow manual correction
- 60 - Continuous retry under operator control
- 43 Phase I Area Indicator. Punch a one-digit code indicating the area allocation procedure to be used by Phase I:
1 - Phase I will decide
2 - Two-area system
3 - Three-area system
- 44-47 Phase I Storage Limit. Four-digit address of the highest 7070 storage location usable by Phase I.
- 48-51 Phase II Storage Limit. Four-digit address of the highest 7070 storage location usable by Phase II.
- 52-55 Phase III Storage Limit. Four-digit address of the highest 7070 storage location usable by Phase III.
- 56-60 Form 3 Length Field. For record Forms 1 and 2, punch five zeros. For record Form 3, punch the position within the record of the field which contains the record word count: columns 56-58 contain the word number, assuming that the first word of the record is numbered 000, and columns 59-60 contain field control.
- 61 Input File Label Indicator. Punch a one-digit code describing the labels of the tape reels making up the input file:
0 - No labels
1 - Low density
2 - High density
- If a 1 or 2 is punched, an Input File Label Control Card (Control Card 3) must be furnished.
- 62 Work Tape Label Indicator. Punch a one-digit code describing the labels on work tapes:
0 - No labels
1 - Low density
2 - High density
- 63 Output File Label Indicator. Punch a one-digit code describing the labels of and for the tape reels to be used for containing the sorted output file:
0 - No labels
1 - Low density
2 - High density
- If a 1 or 2 is punched, an Output File Label Control Card (Control Card 4) must be furnished.

<u>Columns</u>	<u>Contents</u>
64	<u>SPOOL Indicator.</u> Punch the digit 0 if SPOOL programs will not operate concurrently with Sort 90, and the digit 1 if they will or may. If a 1 is punched, electronic switches 29 and 30 must be in the OFF condition during any portion of the sort program in which SPOOL programs are not in operation.
65	<u>Input Checkpoint Indicator.</u> If checkpoints are contained on any of the reels making up the input file, punch 1. Otherwise, punch 0.
66-69	<u>Phase I Edit Indicators.</u> If Form 1 or Form 2 records are being sorted and Phase I editing modifies the record type and/or record length, punch in column 66 the record form and in columns 67-69 the record length (or maximum record length) in 7070 words of a Phase I output record. Otherwise, punch four zeros.
70-73	<u>Phase III Edit Indicators.</u> If Form 1 or Form 2 records are being sorted and Phase III editing modifies the record type and/or record length, punch in column 70 the record form and in columns 71-73 the record length (or maximum record length) in 7070 words of a Phase III output record. Otherwise, punch four zeros.
74-75	Unused by Sort 90 and may be punched as desired. Each column must be filled.
76-80	Zeros.

Control Card 2

<u>Columns</u>	<u>Contents</u>
1-6	<u>Control Data Field 1 Position.</u> Position in the record of the first (major) control data field. This is punched in the form WWWDLL, where WWW is the number of the word within which the control data field begins, assuming that the first word of the record is numbered 000; D is the starting digit within WWW of the control data field, counting the digits from left to right 0 to 9; and LL is the total number of digits contained in the control data field. Nine succeeding six-column fields are used to describe further control data field positions. These must be punched consecutively from left to right. Unused fields must be filled with zeros.
7-12	<u>Control Data Field 2 Position.</u>
13-18	<u>Control Data Field 3 Position.</u>

<u>Columns</u>	<u>Contents</u>
19-24	<u>Control Data Field 4 Position.</u>
25-30	<u>Control Data Field 5 Position.</u>
31-36	<u>Control Data Field 6 Position.</u>
37-42	<u>Control Data Field 7 Position.</u>
43-48	<u>Control Data Field 8 Position.</u>
49-54	<u>Control Data Field 9 Position.</u>
55-60	<u>Control Data Field 10 Position.</u>
61-63	<u>Total Control Data Digits.</u> Punch the total number of digits of control data for the sort, i. e. , the sum of the LLs in columns 1-60.
64	<u>Collating Sequence.</u> Punch a one-digit code specifying the collating sequence to be used by the sort: 0 - Algebraic, ascending 1 - Algebraic, descending 2 - Absolute, ascending 3 - Absolute, descending
65	<u>Summarizing Equals Control.</u> If summarizing in Phase II and/or Phase III and not all control data fields are being considered in determining that two records are identical, punch the number of control data fields (1-9) which are to be considered. If summarizing is not to be performed or if all control fields are to be considered in determining that two records are identical, punch a zero.
66-67	<u>Phase II Summarizing Pass Control.</u> If summarizing in Phase II, punch the number of Phase II passes in which summarizing is to be performed in every Phase II pass, punch 99. Note that these columns must be non-zero if summarizing is to be performed in any Phase II pass.
68-75	Unused by Sort 90 and may be punched as desired. Each column must be filled.
76-80	Zeros.

Control Card 3 (Input File Label Control Card)

Control Card 3 must be furnished if column 61 of Control Card 1 is non-zero. It may be punched with parameters which will be inserted by Sort 90 in the label information entry for the sort input file, and which IOCS will then use to check the labels of the reels of the input file. No parameter will be stored, however, if columns used for it (see below) are punched with zeros.

<u>Columns</u>	<u>Contents</u>
1-7	Unused by Sort 90 and may be punched as desired. Each column must be filled.
8-10	<u>Reel Sequence Number</u> of the first reel of the input file.
11-15	<u>Creation Date</u> of the input file.
16-20	<u>File Serial</u> of the input file.
21-40	<u>File Identification</u> of the input file. Punched in double-digit form.
41-75	Unused by Sort 90 and may be punched as desired. Each column must be filled.
76-80	Zeros.

Control Card 4 (Output File Label Control Card)

Control Card 4 must be furnished if column 63 of Control Card 1 is non-zero. It may be punched with parameters which will be inserted by Sort 90 in the label information entry for the sort output file, and which IOCS will then insert in the label to be written on reels of the output file. No parameter will be stored, however, if the columns used for it (see below) are punched with zeros.

<u>Columns</u>	<u>Contents</u>
1-7	Unused by Sort 90 and may be punched as desired. Each column must be filled.
8-10	<u>Retention Cycle</u> for the output file.
11-30	<u>File Identification</u> for the output file. Punched in double-digit form.
31-74	<u>Comments</u> for the output file label. Punched in double-digit form. If this field is non-zero, its 22 characters will be stored in the 22 high-order positions of the comments field.
75	Unused by Sort 90 and may be punched as desired. The column must be filled.
76-80	Zeros.

MESSAGES

Messages that may be typed out in any run of Sort 90 are listed below approximately in the order in which they might appear. The text of the message is shown in block letters; an explanation follows directly below.

Messages are either procedural, indicating the progress of the program; cautionary, pointing out non-standard procedures; or statistical, furnishing information (counts, etc.) about the file being sorted.

Messages typed by the Input/Output Control System are not included in this listing.

PHASE I

7070 SORT 90

Identifies the Sort 90 program.

NOTE ONE CHAN X

Control card punching indicates that Channel A and Channel B tape units will be attached to the same channel, the number of which is typed in the message. A two-area system of reading, writing and processing will be used in Phase I.

NOTE # INPUT REELS LESS THAN # INPUT DRIVES

The number punched in the Input Reel Count field of Control Card 1 is less than the number of tape units specified for reading the reels of the Phase I input file. While this situation is permissible, an error may have been made in control card punching. If so, depress STOP, correct Control Card 1 and reload Phase I.

3 AREA

Phase I will use a three-area system of reading, writing and processing.

2 AREA G

Phase I will use a two-area system of reading, writing and processing, because the amount of storage available is insufficient to permit the use of a three-area system.

2 AREA T

Phase I will use a two-area system of reading, writing and processing, because its processing time will exceed the sum of the input and output tape times.

BZZZZGXXXX

If Form 1 or 2 records are being sorted, the number of records composing a Phase I sorted output sequence will be XXXX, and the sort blocking factor will be ZZZZ. If Form 3 records are being sorted, the number of records composing a Phase I sorted output sequence will be that number which can fit in XXXX 7070 words, and the maximum sort block size will be ZZZZ words.

The Phase I main program is about to begin. The Phase I assignment program has found no errors in control card punching.

REEL XX DROPPED

The XXth input reel has been dropped from the sort.

REEL XX LAST REEL

The XXth input reel is the last input reel processed by Phase I.

IN	+xxxxxxxxxx
DELET	+xxxxxxxxxx
DUMP	+xxxxxxxxxx
OUT	+xxxxxxxxxx

One, two, three or all of these messages may appear at the end of Phase I. The IN message shows the number of input data records successfully read by Phase I. The DELET message, if typed, shows the number of records deleted. The DUMP message, if typed, shows the number of input file tape blocks written on an unreadable records tape. The OUT message shows the number of records written by Phase I; it will be typed only if records have been deleted.

REPLACE XX

Tape unit XX, from which the last reel of the input file was read, will be used as a work tape unit in Phase II. Dismount the last input reel and mount a work tape on unit XX, unless the input reel is to be destroyed by Phase II. Halt 2111 will occur at the beginning of Phase II to ensure enough time for changing the reel.

PHASE II

PHASE 2 PASSXX REQZZ

The XXth Phase 2 pass is about to begin and ZZ Phase II passes (including the XXth pass) remain to be performed.

CHPT +CUxxxxxNNN

Checkpoint record NNN is about to be written on the tape mounted on tape channel and unit CU.


```

IN          +xxxxxxxxxxx
SUMM       +xxxxxxxxxxx
DELET      +xxxxxxxxxxx
DUMP       +xxxxxxxxxxx, NO RECONCILIATION
OUT        +xxxxxxxxxxx

```

These messages may or may not appear at the end of a Phase II pass. The IN and OUT messages show the number of records read and the number written, respectively, during the pass; these two messages will be typed only if at least one of the other three is typed. The SUMM and DELET messages, if typed, show the number of records summarized and deleted, respectively, during the pass. The DUMP message, if typed, shows the number of input tape blocks written on the unreadable records tape during the pass; it also indicates that the pass-to-pass record count check and pass-to-pass hash total check (if elected) have been suspended during the pass.

PHASE III

PHASE 3

Phase III is about to begin.

```
CHPT +CUxxxxxNNN
```

Checkpoint record NNN is about to be written on the tape mounted on tape channel and unit CU.

```

IN          +xxxxxxxxxxx
INSRT       +xxxxxxxxxxx
SUMM       +xxxxxxxxxxx
DELET      +xxxxxxxxxxx
DUMP       +xxxxxxxxxxx, NO RECONCILIATION
OUT        +xxxxxxxxxxx
7070 SORT 90 ENDS

```

The last message identifies the conclusion of Sort 90 and precedes Halt 3333. The IN message shows the number of records read by Phase III; it is typed only if at least one of the four messages directly below is typed. The INSRT, SUMM and DELET messages, if typed, show the number of records inserted, summarized and deleted, respectively, during Phase III. The DUMP message shows the number of input tape blocks written on the unreadable-records tape during Phase III; it also indicates that the pass-to-pass record count check and pass-to-pass hash total check (if elected) have been suspended during Phase III. The OUT message shows the number of records written by Phase III; it is always typed.

OUTPUT REEL XX TPZZ

The writing of the XXth reel of the output file, which is mounted on tape unit ZZ, has been completed. Sort 90 will now either continue writing records of the output file on another available tape unit, or come to Halt 3115, or conclude.

PROGRAMMED HALTS

All Sort 90 programmed halts are Halt and Proceed (HP) instructions. The halts listed below are tabulated against digits 6-9 of the program register typeout. Thus the typeout

2649 -0000001101

refers to Halt 1101.

For most halts, the sixth digit is 1, 2 or 3, indicating the Sort 90 phase during which the halt occurs. An exception is Halt 7070, which is common to all three phases and indicates a previous typing error.

The letter C or U, or neither, may appear to the right of the halt number in the listing. C means that the halt is conditional, i. e., options are provided for continuing Sort 90; U means that the halt is unconditional, i. e., Sort 90 may not be continued; no letter means that the halt is merely procedural. The text of the message, which may be typed before a halt, appears in block letters to the right of the halt number, immediately above the explanation of the halt.

Programmed halts executed by the Input/Output Control System are not included in this listing.

PHASE I

1000 C CONTROL CARD #X READ ERROR

A validity error was detected while Control Card #X (1, 2, 3, or 4) was being read. If using a card reader, correct the card in error, replace it and the cards behind it (if any) and press START. If the control cards are being read from tape, press START to reread the control card image.

CONTROL CARD #X EOF

An end-of-file signal occurred while Control Card #X was being read. Same action as above.

TAPE CONTROL CARD #X FSW CODE #Y

A condition other than Correct Length Record, Error or End of File was detected while reading Control Card #X from tape. #Y will be the condition code digit of the final status word. Press START to reread the control card image.

1001 U CONTROL CARD ERROR OF TYPE #XX

From its analysis of control card punching, the Phase I assignment routine has discovered a situation which makes it impossible for the sort to continue. The type of error is

indicated by the code number XX in the list below. In all cases, the reason for the halt is obvious. To resume the sort, the control card in error must be repunched correctly and both Phase I and the control cards reloaded.

- 01 Two or more of the tape units to be used for Phase I input have the same number.
- 02 Two or more of the Channel A tape units have the same number.
- 03 Two or more of the Channel B tape units have the same number.
- 04 Less than two tape units are indicated for Channel A and/or Channel B.
- 05 Although Channel A and Channel B differ, Channel B and the Input Channel are the same.
- 06 A Channel A tape unit is the same as a Channel B tape unit.
- 07 A tape unit to be used for reading the input file is the same as a Channel B tape unit.
- 08 The tape unit to be used for dumping unreadable records is the same as an input file tape unit or a work tape unit.
- 09 More than 160 digits of control data information have been specified.
- 10 The number punched in the Summarizing Equals Control column is greater than the total number of control data fields specified.
- 11 One or more of the control data fields falls beyond the indicated minimum length record.
- 12 The total number of digits in the specified control data fields does not agree with the number punched in the Total Control Data Digits columns.
- 13 The specified control data fields include more than 16 control data segments.
- 14 The fifth and sixth positions (LL) of one or more otherwise valid control data field descriptions are punched with zeros, indicating a control data field(s) of no length.
- 15 The number punched in the Output Tape Unit Type column is other than 1, 2, 3, or 4.
- 16 The number punched in the Record Form column is other than 1, 2 or 3.

- 17 The number punched in the Record Count Indicator column is greater than 3.
- 18 The number punched in the Hash Total Indicator column is either a zero, although there is at least one non-zero punch in the Hash Total Field columns, or greater than 7.
- 19 The number punched in the record type column of the Phase I Edit Indicators field is greater than 2, or indicates an unpermitted mode of record form alteration.
- 20 The number punched in the record type column of the Phase III Edit Indicators field is greater than 2, or indicates an unpermitted mode of record form alteration.
- 21 The number punched in the Collating Sequence column is greater than 3.
- 22 The record-type column of the Phase I or Phase III Edit Indicators field contains a valid non-zero punch, but the Phase I length field or Phase III length field, respectively, contain zeros, indicating that Phase I or Phase III editing will create records of no length.
- 23 The number punched in the Input Channel, Channel A or Channel B columns is other than 1 or 2.
- 24 The number punched in the Checkpoint Indicator column is neither 0 nor 1.
- 25 The Hash Total Field is located beyond the last word of any or all applicable record length descriptions.
- 26 The Form 3 Length Field is located beyond the last word of the indicated shortest Form 3 record.
- 27 The number punched in the Phase I Area Indicator column is either zero or greater than 3.
- 28 The Input Record Length field is punched with zeros.
- 29 A control data field description is punched with six zeros, although a control data field description to the right is non-zero.
- 30 The number punched in the Zero Suppression Indicator column is greater than 1.
- 31 The number punched in the Input Tape Type column is either zero or greater than 4.
- 32 The Input Blocking field is punched with zeros.

- 33 The Output Blocking field is punched with zeros.
- 34 The number punched in the Unreadable Record Indicator field is other than 00, 10-29, 50 or 60.
- 35 Input tapes are indicated as unlabeled, but a hash total and/or a record count option asks Phase I to check fields in input tape trailer labels.
- 36 Input tapes are indicated as unlabeled, and the Input Reel Count columns are punched with zeros.
- 37 Output tapes are indicated as unlabeled, but a hash total and/or a record count option asks Phase III to insert fields in output tape trailer labels.
- 38 The number punched in the Input File Label Indicator column is greater than 2; or a low density input file with high density labels is specified.
- 39 The number punched in the Output File Label Indicator column is greater than 2; or a low density output file with high density labels is specified.
- 40 The number punched in the Work Tape Label Indicator column is greater than 2.
- 41 This code applies to Form 3 records only. The indicated minimum input record size is greater than the indicated maximum input block size.
- 42 The number punched in the SPOOL Indicator column is neither 0 nor 1.

1002 C UNEQ NUMBER A, B TPS

The number of tape units indicated for Channel A is unequal to the number indicated for Channel B. If desired, repunch Control Card 1 and reload Phase I. If the control card is not corrected and the Start key is depressed, the order of merge to be used in Phases II and III will be adjusted to the lower of the two numbers; the excess tape unit(s) on the channel containing the greater number will not be used.

1003 C INPUT CHAN X NOT WORK CHAN

Channel A and Channel B tape units are specified as using the same channel, but the input file is to use a different channel. If only one tape channel is available, this is a control card punching error; correct Control Card 1 and reload Phase I. If two tape channels are available, the sort will continue if the Start key is depressed; however, the indicated configuration will not permit overlapping of tape reading and writing in Phases II and III.

1004 C NOTE HASH FIELD WORD 000 DIG 0

The Hash Total Indicator column is punched with a valid hash total option, but the Hash Total Field columns are punched with zeros. If it is intended that hash totals be taken on the first digit of the first word of each record, press START. If not, correct Control Card 1 and reload Phase I.

1005 C CONTROL FLD OVERLAP

XX,XX

One or more pairs of control data fields are overlapped in a record. The first line will be typed once. The second line will be typed as many times as there are overlapping pairs of control data fields, where XX,XX will be the numbers of the overlapping fields. If punched improperly, correct Control Card 2 and reload Phase I. If the overlaps are desired, press START.

1006 C LENGTH FIELD QUESTIONABLE

The Form 3 Length Field contains enough digit positions to identify a record longer than the indicated longest input or output tape block. If improperly punched, correct Control Card 1 and reload Phase 1. If the Form 3 Length Field contains high-order zeros, and if it is known that no single record will exceed in length either of the indicated maximum block sizes, press START.

1010 C An indicator word in 7070 storage suggests that Sort 90 has been previously interrupted and restarted and that Phase III should now be loaded into storage. If this is the case, set Alteration Switch 1 OFF and press START. Halt 1011 will follow immediately. If, however, Sort 90 is being started from the beginning, set Alteration Switch 1 ON and press START. Note that the halt will often occur if Sort 90 is being started from the beginning after having been previously discontinued in Phase II.

1011 C Sort 90 will continue to load itself until Phase III has been loaded and can execute itself. At this point it is necessary to know if the Sort 90 control cards and the sort program are being read through the same input device. If so, set Alteration Switch 1 ON and press START; the control cards will be bypassed. If not, set Alteration Switch 1 OFF and press START.

1012 U A card read as a Sort 90 control card does not contain zeros in columns 76-80. It is likely that an insufficient number of control cards were furnished and that a Phase II program card has been read. Correct the situation and reload Phase I.

1051 U NOT ENOUGH STORAGE AVAILABLE PHASE X

Insufficient storage is available for the indicated phase. Check the Storage Limit address for the phase and, also, the record length and input and output blocking descriptions.

1052 C 3AREA G OF XXX

Insufficient storage is available for Phase I to use a three-area system of reading, writing and processing. For record Forms 1 and 2, XXX will be 000, indicating that no input tape blocks can be read. For record Form 3, XXX will be the maximum number of words which Phase I can permit in an input tape block, the number being smaller than the maximum tape block size punched in Control Card 1. If the Start key is depressed, the program will attempt to use a two-area system of reading, writing and processing.

1053 U 2AREA G OF XXX

Insufficient storage is available for Phase I to use a two-area system of reading, writing and processing. For record Forms 1 and 2, XXX will be 000, indicating that no tape blocks can be read. For record Form 3, XXX will be the maximum number of words which Phase I can permit in an input tape block, the number being smaller than the maximum tape block size punched in Control Card 1. Check the Storage Limit address for Phase I and, also, the record length and input and output blocking descriptions.

1054 C MAX BLOCK XXXX

This halt pertains only to Form 3 records. The number of words in the computed sort tape block, XXXX, is less than the indicated maximum number of words composing an input and/or output tape block. If Control Card 1 punching is correct and it is known that no single data record will exceed in length the computed sort block size, press START.

1101 C RECORD CNT DISCREPNCY +XXXXXXXXXX +XXXXXXXXXX
HASH TOTL DISCREPNCY +XXXXXXXXXX +XXXXXXXXXX

Either or both of the above messages may be typed before the halt. Each indicates a discrepancy between a value found in the trailer label of the current input reel and a value computed by Phase I. The first numerical word will be the label value, and the second will be the computed value. There may also be typed by itself or with one or both of these messages an IOCS message noting a discrepancy between the block count accumulated in Phase I and that contained in the trailer label of the current input reel. To ignore the discrepancies and continue Phase I processing, set Alteration Switch 1 ON and press START. To attempt remedial action, set Alteration

Switch 1 OFF and press START; Halt 1102 will follow immediately.

1102 C Two remedial actions are possible: To drop all records of the current input reel from the sort, set Alteration Switch 1 OFF and press START. To reprocess the current input reel, set Alteration Switch 1 ON and press START.

1103 U In performing a Segment Mark Backspace per Count operation, the load point was reached before the desired segment mark was detected. This situation is beyond the control of Sort 90. Discontinue and restart the sort from the beginning.

1104 C LABEL EOR, COUNT 00
LABEL EOF, COUNT XX

Either the trailer label of the current input reel or the Input Reel Count, but not both, indicates that the current reel is the last input reel. The first message will be typed if the Input Reel Count, but not the trailer label, indicates that the current input reel is the last. The second message will be typed if the trailer label, but not the Input Reel Count, indicates that the reel is the last; XX will be the number of input reels still to be read according to the Input Reel Count. In either case, set Alteration Switch 1 OFF and press START to enter the Phase I end-of-job routine, accepting the current reel as the last; or set Alteration Switch 1 ON and press START, to continue processing reels of the input file.

1105 U In performing a Segment Mark Backspace per Count operation, a tape mark was reached before the desired segment mark was detected. This situation is beyond the control of Sort 90. Discontinue and restart the sort from the beginning.

1107 C The current Phase I input reel has now been completely processed and Halt 1110 or 1114 has previously occurred. To terminate Phase I at this point so that the current input reel will be considered the last input reel, whether or not it is, set Alteration Switch 1 OFF and press START. To continue Phase I, permitting the program to determine whether or not the reel is the last, set Alteration Switch 1 ON and press START.

1108 C A validity error has occurred during a Write Segment Mark operation. To perform a Tape Skip and rewrite the segment mark, press START.

1109 U RECORD TOO SMALL

The length field of a Form 3 record indicates that the record is shorter than the minimum length punched in the Input Record Length field of Control Card 1. If the Start key is depressed, the offending record will be typed. The sort cannot continue. Correct Control Card 1 and restart the sort from the beginning.

- 1110 C All but one of the Phase I output tapes have been completely filled with records. The total number of records being sorted has therefore just reached or exceeds the maximum permissible number. To drop the records of the current input reel from the sort and enter the Phase I end-of-job routine, set Alteration Switch 1 OFF and press START. If it is known that few input records remain, that one or more of the Phase I output tapes is short, or that heavy summarizing will be performed in an early Phase II pass, Phase I may be continued by setting Alteration Switch 1 ON and depressing the Start key.
- 1112 C All of the Phase I output tapes have been completely filled. To drop the records of the current input reel from the sort and enter the Phase I end-of-job routine, press START. Providing Halt 1110 or Halt 1114 has not occurred previously, the sort will go to a successful completion. If either Halt 1110 or Halt 1114 has occurred previously, the number of records dangerously exceeds the allowable maximum. The sort will most likely not succeed unless one or more of the Phase I output tapes is short, or unless heavy summarizing will be performed in an early Phase II pass.
- 1114 C Although more than the maximum allowable number of records have been written on the Phase I output tapes, if Phase I is terminated at this point the sort will proceed to a successful completion. To terminate Phase I, set Alteration Switch 1 OFF and press START; the records of the current input reel will be dropped from the sort and Phase II will begin. If instead it is desired to continue processing input records, set Alteration Switch 1 ON and press START.
- 1124 C A Long Length Record has been detected in the sort input file. This probably means incorrect control card preparation. If the Start key is depressed, as much of the record as was entered into 7070 storage will be typed, the record will be discarded and Phase I will continue.
- 1126 C A segment mark has been detected in the sort input file. If the Start key is depressed, the segment mark will be ignored and Phase I will continue.
- 1127 C A Short Character Length Record has been detected in the sort input file. This probably means that the wrong input file has been mounted. If the Start key is depressed, the record involved will be typed and discarded and Phase I will continue.
- 7070 C An error signal was detected while typing the message which immediately precedes the Halt 7070 Program Register typeout. To repeat the typing, set Alteration Switch 1 ON and press START. To ignore the error, set Alteration Switch 1 OFF and press START. Since most messages are followed by a programmed halt, be prepared for another (and different) halt.

PHASE II

2000 C SEQUENCES +xxxxxxxxxxx +xxxxxxxxxxx

The first numerical word will be the number of sorted sequences produced by the previous Phase II pass (or by Phase I if the Phase II pass just concluded is the first); the second numerical word will be the number of sorted sequences produced by the Phase II pass just concluded. The second number is not less than the first, implying that the Phase II pass just concluded did not further the sorting process. To continue with the next Phase II pass, press START. This action is not recommended, however, because the next Phase II pass is not likely to succeed in reducing the number of sequences.

2001 C PHASE 2 PASSXX REQZZ

The number of Phase II passes, ZZ, that are still required at the beginning of pass XX is not less than the number required at the beginning of the previous pass. However, the previous pass succeeded in reducing the number of sorted sequences, indicating that some progress in the sort is being made. To continue with pass XX, press START. This action is recommended.

2002 C Modifications made in accordance with Modification II-4 indicate that summarizing is to be performed during the next Phase II pass; however, the Summarizing Pass Control column of Control Card 2 was punched with zeros and, therefore, the special comparison routine required in Phase II for summarizing has not been generated. Correct Control Card 2 and restart the sort from the beginning or press START to continue Phase II without summarizing.

2003 An indicator word in 7070 storage suggests that Sort 90 has been previously interrupted and restarted and Phase III should now be loaded into storage. (Halts 1010 and 1011 may have just occurred.) If so, set Alteration Switch 1 OFF and press START. If, however, it is known that Phase II is to be executed at this time, set Alteration Switch 1 ON and press START.

2004 U Form 3 records are being sorted; however, the first instruction of the area used in Phase II for Form 3 assignment routine programming is not as expected. The cards containing such programming have probably been removed from the Phase II section of the Sort 90 program deck. Put the missing cards in the deck and restart the sort from the beginning.

2005 C An indicator word in 7070 storage suggests an improper action in loading Phase II. The contents of storage may have been disturbed between Phases I and II (for example, by improperly prefacing Phase II by a zero storage routine) or an incorrect checkpoint record may have been read into

storage. The offending condition should be removed and the sort restarted from the beginning. Depressing the Start key will cause the condition to be ignored and Phase II to continue; the sort may not continue correctly, however, if sections of storage other than the indicator word location have been altered.

- 2101 U In concluding the processing of an input tape reel, the IOCS end-of-reel exit has been used. Because all sort work tapes contain end-of-file trailer labels, this halt should never occur. It suggests either faulty tape handling or improper modification of IOCS routines. If the halt occurs, restart the sort from the beginning.
- 2103 U All of the output tapes for the current Phase II pass have been filled. Either one or more of the tape reels now being used for output are too short, or the number of records being sorted exceeds the maximum permissible number. If the number of records being sorted is too great, restart the sort from the beginning with a smaller number of input records.
- 2104 C All but one of the output tapes for the current Phase II pass have been filled. Either one or more of the output reels is too short, or the number of records being sorted exceeds the maximum permissible number. To continue the current pass, press START. If the remaining output tape cannot contain the remaining input records, Halt 2103 will later occur. Under certain conditions the pass may be concluded properly, but results in Halt 2000 or 2001 at the beginning of the next Phase II pass.
- 2105 U In performing a Segment Mark Forwardspace per Count operation, a tape mark was reached before the desired segment mark was detected. This situation is beyond the control of Sort 90. Discontinue and restart the sort from the beginning.
- 2108 C A validity error has occurred during a Write Segment Mark operation. To perform a Tape Skip and rewrite the segment mark, press START.
- 2111 C This halt will occur at the beginning of Phase II to provide time for replacing the last reel of the input file by a work tape. When the tape unit described in the message REPLACE XX (see page 35) has been mounted with a work tape, or if the last input reel is to be destroyed by Phase II, press START.
- 2199 C IN +xxxxxxxxxx
SUMM +xxxxxxxxxx
DELET +xxxxxxxxxx
OUT +xxxxxxxxxx
RCRD CNT DISCREPNCY
HASH TOTL DISCREPNCY +xxxxxxxxxx +xxxxxxxxxx

The pass-to-pass record count and/or pass-to-pass hash total

check (if elected) are off in the current Phase II pass. The IN message shows the number of records written by the previous pass (or by Phase I if the current Phase II pass is the first). The SUMM and DELET messages, if typed, show the number of records summarized and deleted, respectively, during the current pass. The OUT message shows the number of records written by the current pass. The first numerical word in the HASH TOTL DISCREPNCY message gives the hash total for records written by the previous Phase II pass (or by Phase I if the current pass is the first); the second numerical word gives the sum of the hash totals of records written, deleted and summarized during the current pass. If record processing routines have been added to Phase I and/or Phase II, this halt generally indicates incorrect coding of the routines. To ignore the discrepancy and enter the next Phase II pass, press START. This action is not recommended.

7070 C An error signal was detected while typing the message which immediately precedes the Halt 7070 Program Register typeout. To repeat the typing set Alteration Switch 1 ON and press START. To ignore the error, set Alteration Switch 1 OFF and press START. Since most messages are followed by a programmed halt, be prepared for another (and different) halt.

PHASE III

3004 U Form 3 records are being sorted; however, the first instruction of the area used in Phase III for Form 3 assignment routine programming is not as expected. The cards containing such programming have probably been removed from the Phase III section of the Sort 90 program deck. Put the missing cards in the deck and restart the sort from the beginning.

3005 C An indicator word in 7070 storage suggests an improper action in loading Phase III. It may be that the contents of storage have been disturbed between Phases II and III (for example, by improperly prefacing Phase III with a zero storage program); or that an incorrect checkpoint record has been read into storage. The offending condition should be removed and the sort restarted from the beginning. Depressing the Start key will cause the condition to be ignored and Phase III to continue; the sort may not continue correctly, however, if sections of storage other than the indicator word location have been altered.

3101 U In concluding the processing of an input reel, the IOCS end-of-reel exit has been used. Because all sort work tapes contain end-of-file trailer labels, this halt should never occur. It suggests either improper modification of IOCS routines or faulty tape handling. If the halt occurs, restart the sort from the beginning.

3105 U In performing a Segment Mark Forwardspace per Count operation, a tape mark was reached before the desired segment mark was detected. This situation is beyond the control of Sort 90. Discontinue and restart the sort from the beginning.

3109 C SEQUENCE BREAK

A sequence break has been detected between the current sort output record and the preceding sort output record. This may result from improper Phase III modifications which insert records into the sort output file (See Modification III-2, page 66), or from manual correction of an input record at some point in the sort in which the control data was involved. To type the current sort output record and the preceding output record, set Alteration Switch 1 ON and press START; Halt 3110 will occur after the typing is complete. To ignore the sequence break, set Alteration Switch 1 OFF and press START; this will cause the current output record to be included in the sort output file and the next sequence-check comparison to be made between the current record and the next output record. This action will result in an out-of-sequence condition in the sort output file; in recognition, the message SEQUENCE BREAK TO OUTPUT will be typed.

3110 C This halt allows for optional treatment of the sequence break condition discovered at Halt 3109. To drop the current output record from the sort output file, set Alteration Switch 1 ON and press START. This will cause the next sequence-check comparison to be made between the preceding output record and the next output record; also, the message DROP will be typed. To ignore the sequence break, set Alteration Switch 1 OFF and press START. This will cause the current output record to be included in the sort output file and the next sequence-check comparison to be made between the current record and the next output record. This action will result in out-of-sequence condition in the sort output file; in recognition, the message SEQUENCE BREAK TO OUTPUT will be typed.

Note that the condition responsible for the sequence break should be investigated and corrected, even though the two options are provided.

3111 C This halt will occur at the beginning of Phase III, if no Phase II passes have been required, to provide time for replacing the last reel of the input file by a work tape. When the tape unit described in the message REPLACE XX (see page 35) has been mounted with a work tape, or if the last input reel is to be destroyed by Phase III, START.

3115 C Full reels of output records have been written on all tape units available for the sort output file, as indicated by previous OUTPUT REEL XX TPZZ messages. These same tape units will be used for further reels of the output file. Remove the

completed reels of the output file and mount work tapes on these tape units, and then press START.

3199 C IN +xxxxxxxxxxx
INSRT +xxxxxxxxxxx
SUMM +xxxxxxxxxxx
DELET +xxxxxxxxxxx
OUT +xxxxxxxxxxx
RCORD CNT DISCREPNY
HASH TOTL DISCREPNY +xxxxxxxxxxx +xxxxxxxxxxx

The pass-to-pass record count and/or the pass-to-pass hash total check (if elected) are off. The IN message shows the number of records written by the last Phase II pass (or by Phase I if Phase II was not required). The INSRT, SUMM and DELET messages, if typed, show the number of records inserted, summarized and deleted, respectively, during Phase III. The OUT message shows the number of records written by Phase III. The first numerical word in the HASH TOTL DISCREPNY message gives hash total for records written by the last Phase II pass (or by Phase I if Phase II was not required); the second numerical word gives the sum of the hash totals of records written, deleted and summarized during Phase III. If record processing routines have been added to any phase of Sort 90, this halt generally indicates incorrect coding of the routines. To ignore the discrepancy and enter the Phase III end-of-job routine, press START. This action is not recommended.

3333 7070 SORT 90 ENDS

If the Start key is depressed, the program will enter an IOCS end-of-job "busy loop" to permit the continued execution of SPOOL programs.

7070 C An error signal was detected while typing the message which immediately precedes the Halt 7070 Program Register typeout. To repeat the typing, set Alteration Switch 1 ON and press START. To ignore the error, set Alteration Switch 1 OFF and press START. Since most messages are followed by a programmed halt, be prepared for another (and different) halt.

MODIFICATIONS

The Sort 90 program offers a convenient opportunity for performing actions upon the file of records being sorted; these actions may lead directly out of the sorting process, such as summarizing, or may simply take advantage of the several tape passes which the file undergoes during the sort. This section of the Sort 90 manual consists of two parts: First, a general description of the procedure to be used in modifying the program; and second, a series of detailed instructions for effecting several common modifications. In making these modifications, a detailed knowledge of the Sort 90 program is not required.

In detailing suggested modifications, symbolic labels are used when referring to instructions and constants within Sort 90. Accompanying the Sort 90 program deck is a listing of 7070 addresses corresponding to each of the symbolic labels, together with a description of the storage areas available for modifications and the numbers of available index words and electronic switches. It is suggested that modifications be coded in 7070 Autocoder language, using statements such as EQU, ORIGIN CNTRL, XRESERVE CNTRL, and SRESERVE CNTRL to cause the assembled modification to fall into the desired storage areas and make correct reference to locations within Sort 90.

In performing the first few runs of Sort 90 after it has been modified, it is recommended that full use be made of the pass-to-pass hash total checking features available in Sort 90. Hash total checks are particularly important if the modifications deal with processing the records of the file being sorted, as they often detect mishandling of the records.

PATCHING

The suggested technique for modifying Sort 90 is the use of "patches," or condensed load cards which contain the instructions required to effect the modification and which are inserted immediately before the execute card of the Sort 90 phase which they are to modify. Thus the loading sequence consists of the Sort 90 phase, followed by new and altered instructions, followed by the execute card which initiates execution of the program.

It is also possible to alter the contents of Sort 90 condensed load cards directly, but this technique is to be frowned upon as it does not leave intact a record of what instructions originally composed the program.

TYPES OF MODIFICATIONS

One type of modification alters the program logic of Sort 90, usually through the insertion of revised or new program steps. An instruction in the Sort 90 program, called an "exit" instruction, is overlaid through "patching" with a Branch instruction to a modification routine. The last-executed instruction of the modification routine is a Branch to a "return" instruction within Sort 90. In some cases the modification routine will conditionally branch to one of several possible return instructions. In a few cases, no new routines are needed, and an exit instruction will be overlaid with a


Branch directly to a return instruction, thus bypassing a part of the Sort 90 program, or realigning the order in which sections of it are executed.

Another type of modification simply alters a constant or table entry within Sort 90. Modification I-4 details a number of possible modifications of this type.


PLACEMENT OF MODIFICATIONS

Each Sort 90 phase consists of two sections: a main program, or "running program," which always remains in 7070 storage and actually effects the sorting of records; and an assignment routine which is executed prior to entering the main program and which modifies and adjusts the main program for the desired Sort 90 application. The assignment routine is destroyed by the generation of RDWs and the read-in of records after it has relinquished control to the main program. Phase I loads in two sections: the first section consists of the main program, other than IOCS subroutines, and most of the assignment routine; and the second section consists of the IOCS subroutines and the remaining portion of the assignment routine. (The IOCS OPEN subroutine is located, in Phase I, in the second portion of the assignment routine; in Phase II, in the main program; and in Phase III, in the assignment routine.)

Storage allocation for each Sort 90 Phase is as follows:



The point marked "Upper Storage" is controllable by the user for each phase by means of addresses punched in Control Card 1. Modifications to the assignment routine should be placed in the area marked "Not Loaded"; these will be destroyed when the main program is entered. Modifications to the main program may not occupy this area, and must be positioned above the point marked "Upper Storage." In other words, main program modifications should be "crowded" into the uppermost usable section of 7070 storage, and the address of the location immediately below the first modification instruction should be entered on Control Card 1 for the affected phase. If both assignment routine and main program modifications were used in a Sort 90 Phase, area allocation after loading would look as follows:



The phase would use for RDW generation and record read-in the area from the end of the main program up to, but not including, the first word of the main program modifications.

The modifications detailed on the following pages are labeled "Main Program"

or "Assignment Routine," according to which section of the phase they will affect. Phase I assignment routine modifications are labeled "Part 1" or "Part 2." Modifications are serialized by a I, II, or III, to indicate the phase, and a numerical sequence number within each phase.

SETTINGS OF INTERNAL LATCHES

The assignment routine of each phase of Sort 90 executes a Halt Mode Sign Change (HMSC) instruction, thus causing the computer to stop if a sign change should occur; none should occur in correct operation of Sort 90. If a modification is coded which requires the sensing of sign change, it is recommended that a Sense Mode Sign Change (SMSC) instruction be given at the beginning of the modification, and that a HMSC instruction be given before returning to Sort 90.

The assignment routine of each phase of Sort 90 executes a Halt Mode Field Overflow (HMFV) instruction if hash totals are not to be taken during that phase, or a Sense Mode Field Overflow (SMFV) instructions if hash totals are to be taken during that phase. The accumulation of a hash total is the only action of Sort 90 which, in correct operation of the program, might cause a field overflow to occur. The following table identifies the field overflow mode for each Sort 90 phase, depending upon the hash total option punched in column 35 of Control Card 1:

<u>Option</u>	<u>Phase I</u>	<u>Phase II</u>	<u>Phase III</u>
0	Halt	Halt	Halt
1	Sense	Halt	Halt
2	Sense	Sense	Sense
3	Halt	Halt	Sense
4	Sense	Sense	Sense
5	Sense	Halt	Sense
6	Sense	Sense	Sense
7	Sense	Sense	Sense

If a modification is coded which requires the sensing of field overflow during a phase which uses the halt mode, it is recommended that a SMFV instruction be given at the beginning of the modification and that a HMFV instruction be given before returning to Sort 90.

CHECKPOINT AND RESTART

When adding modification routines to the Phase II or Phase III main program, the user should cause the storage area(s) which they occupy to be written with the checkpoint records written during these phases to permit later interruption and restart. Modifications II-6 and III-7 should be consulted in this regard.

EDITING

Sort 90 offers considerable flexibility for addition by the user of routines to edit the records of the file being sorted. Editing routines may be added to Phase I (see Modification I-6) or to Phase III (see Modification III-2), or

to both phases. Following are the changes in record form and length which editing routines may cause in either phase:

If the input form is	The output form may be	And the input data record may be
1	1 or 2	Lengthened or shortened.
2	1 or 2	Lengthened or shortened.
3	3	Shortened, but not lengthened.

Sort 90 will itself take into account any changes made in record length or form through editing. If a record is to be shortened in length, the terminal words of the storage area into which it is read will not be written. If a record is to be lengthened, the area into which the record is read will be followed by unused storage words to permit additional words to be stored and later written with the record.

If Form 2 records are altered in length, the record mark word signalling the end of each record must be repositioned by the editing routine to indicate properly the new length of the record. If Form 3 records are shortened, the contents of the length field of each record must be altered to contain the count of the number of words remaining in the record; the position of the length field itself may not be altered.

SUMMARIZING

By "summarizing" is meant the accumulation of data fields of several records with identical control data and deletion of all but the one record which contains the accumulated data fields of itself and the other records. An example is the summarizing of a file of purchase orders from several departments of a company so as to yield a file of records showing total quantity and price of a given item (the item number being the control data for sorting) to be purchased by all departments collectively. The goal of summarizing is to reduce sharply the number of records of a file while retaining all pertinent information in the file.

Sort 90 offers exit instructions in Phase II and Phase III for branching to added routines for accumulation of data fields of records with identical control data. Whenever the sort executes one of these exit instructions, two records with identical control data are on hand, and the data fields of the second may be added to those of the first. When return is made to the sort, it will delete the second record and test for equality of control data between the next record of the file and the first, or "base" record.

Only by summarizing in Phase III can the user ensure that every "set" of records with identical control data will be reduced to a single record. The comparison routine which tests for identity of control data between two records in Phase III is an integral part of Phase III, and is also used to sequence-check the records written in the sort output file.

Summarizing in Phase II may cause the number of records undergoing further Phase II passes and Phase III to be reduced enough to lower the total running time of the sort significantly. The comparison routine which tests for identity

of control data between records in Phase II is, however, a special routine which is present in Phase II only when the user elects to summarize in one or more Phase II passes; an additional comparison cycle is required for each record on Phase II passes during which summarizing is to be performed.

Procedures for summarizing in Phases II and III are detailed in Modifications II-4 and III-2, respectively.

Control Card 2 allows the user to elect two features which may facilitate his job of summarizing. Columns 66-67 may be punched with a number which will cause Phase II to enter the summarizing mode (i. e. , to perform the added comparison routine necessitated by summarizing) when that number of Phase II passes remains to be executed. Thus if this field were punched 02, summarizing could be performed on the last and next-to-last passes of Phase II. If a number equal to or larger than the number of Phase II passes required is punched, summarizing may be performed on each Phase II pass. In addition, Modification II-5 details procedures which may be followed to gain pass-by-pass control over establishment of the summarizing mode. Regardless of all other considerations, columns 66-67 of Control Card 2 must contain a non-zero number if summarizing is to be performed in any Phase II pass.

The second feature is that column 65 of Control Card 2 may be punched with a number representing a number of control data fields to which Sort 90 will limit itself in determining whether two records have identical control data. (The records will of course be sorted on the full number of control data fields; the punch in column 65 merely controls the use of exit instructions for summarizing.) Thus, if records were to be sorted upon three control data fields but the minor control data field differed among records which were to be summarized, the digit 2 should be punched. Punching 0 in this column is equivalent to punching the full number of control data fields — all control data fields will be considered in identifying records with equal control data. Punching a number greater than the number of control data fields will cause an error halt to occur.

The summarizing routine added to Phase III, or to Phases II and III, need consist only of the instructions which perform the accumulation of data fields of the summarizable records. Deletion of summarized records and reconciliation of record counts and (if any) hash totals are handled by Sort 90 and need not be considered by the user.

PHASE I - MODIFICATIONS

SORT 90 MODIFICATION I-1

Assignment Routine
(Part 1)

Purpose:

To control the unit record synchronizer, or tape channel and unit, from which the Sort 90 control cards are to be read; or to provide for their having been left in 7070 storage by a preceding program.

Description:

1. To read the control cards from a unit record device other than a 7500 Card Reader on Synchronizer 1, overlay location ACCUNITRD with the instruction

UR S, 0+ATEMP

where S is a synchronizer number in the range 1 - 3, or 4 if the 7501 Console Card Reader is to be used.

2. Control cards may be read from a tape unit. If so, it is assumed that the tape containing the control card data is positioned with the read/write head directly before the control card data. Each control card should be one tape record consisting of a mode change character followed by eight numerical words. The density of the tape may be either low or high but the tape unit is assumed already to be set to whichever density is proper.

To cause control card reading from tape, overlay location ACCDUMMY with the instruction

B ACCTAPERTN

and overlay location ACCTAPERD with the instruction

PTR CU, 0+ATEMP

where CU is the channel and unit (10-29) of the tape unit to be read.

3. If the control card data is left in memory before the loading of Phase I, overlay location ACCDUMMY with the instruction

B AASSIGN

The control card data should be left in memory as thirty-two numerical words, the first in location ACNTRLCARD and the last in ACNTRLCARD+31.

Purpose:

Exits for miscellaneous routines at the beginning and ending of Part 1 of the Phase I assignment routine.

Label of Exit
Instruction

Place and Use

AASSIGN	This exit instruction immediately follows the routine which reads the Sort 90 control cards, and may be used for any general assignment routine type of modification. Return should be made to instruction AASSIGN+1.
ASSIGNLINK	This exit instruction occurs at the completion of Part 1 of the Phase I assignment routine. It branches to a routine which clears storage to zero from the end of Part 1 of Phase I to the Phase I Storage Limit (columns 44-47 of Control Card 1). The instruction at ASSIGNLINK+1 is the last instruction of Part 1 of Phase I, and is a Branch to location 0000 to initiate the loading of Part 2. Exit ASSIGNLINK may be used for any general assignment routine type of modification; return should be made to AZERFIRST. If a loading system is used which does not leave a "pivot" to itself in location 0000, ASSIGNLINK+1 should be used to exit to the loading procedure, instructions for which must be located either above the Phase I Storage Limit or in an area of storage below Sort 90.

Note: If after interrupting Phase II and later restarting, Phase I discovers that it has been reloaded, Phase III being desired, Sort 90 will branch directly to ASSIGNLINK+1.

The user should note that at either of the above two exits none of the Input/Output Control System routines have been loaded, and that any attempt to execute IOCS macro-instructions will result in disaster. Part 1 of Phase I assumes itself to be operating with the priority mask set completely to the "disallow" state, and apart from reading control cards performs no input/output actions.

Purpose:

Exits for miscellaneous routines at the end of Part 2 of the Phase I assignment routine.

Label of Exit
Instruction

Placement and Use

AADDASSIGN This is the last instruction of the Phase I assignment routine, and may be used to enter any general housekeeping modification. Return should be made to AGENERATE, following which the generation of RDWs and read-in of records will destroy Part 2 of the Phase I assignment routine. Since the IOCS OPEN subroutine is also destroyed, OPENing of additional tape files in Phase I would be accomplished through this exit.

SORT 90 MODIFICATION I-4

Assignment Routine
(Part 1)

Purpose:

Improvement of the accuracy of the Phase I area allocation computation and selection routine.

Description:

When the Sort 90 Control Card 1 punching indicates that the sort itself is to decide upon using a two-area or three-area read/write system in Phase I, the computation and selection routine is strongly influenced by the values of six 10-digit numerical constants in the Phase I assignment routine. If the computation and selection routine is to be used, and if the standard values do not agree with the actual circumstances of the Sort 90 application, the user should modify these constants.

The symbolic label, meaning, contents, and format of each of the six constants are listed below. The caret sign beneath a constant indicates an implied decimal point; the user should be very careful when altering one of the constants to preserve its decimal-point alignment.

<u>Symbolic Label</u>	<u>Meaning</u>	<u>Format</u>
CACDS	Number of control data segments, beyond the first, which are involved in a comparison between any two typical records of the file being sorted. Assumed zero.	+000000000 ₀ [^]
TAP	Execution time (in milliseconds per record) of added Phase I programming (such as editing) which processes input records. Assumed zero.	+000000000 ₀ [^]
ASL	Average number of data records contained in a natural sequence in the Phase I input file. Assumed two and one-half records; might be much larger if the input file is already partially ordered.	+000000002 ₅ [^]

<u>Symbolic Label</u>	<u>Meaning</u>	<u>Format</u>
AWLI	Average number of tape characters constituting one 7070 word of a Phase I input record. Assumed seven and one-half.	+0000000075 _Λ
AWLS	Average number of tape characters produced by writing for Phase II one word of a data record (zero-suppression will be used). Assumed seven and one-half.	+0000000075 _Λ
AVRL3	(Applies only if Form 3 records are being sorted.) Average length of a Form 3 data record of the sort input file. If this word contains zeros, Sort 90 will assume the average Form 3 data record length to be either the maximum block length or twice the minimum data record size (both quantities come from Control Card 1), whichever is smaller.	+0000000000 _Λ

SORT 90 MODIFICATION I-5

Assignment Routine (Part 1)

Purpose:

To control the size of tape blocks written during Sort 90.

Description:

The Phase I assignment routine computes the sort blocking factor, or block size, to be used during Phases I and II of Sort 90 based only upon storage availability in Phases I, II, and III. It is possible with a two-way merge and 10,000 storage words to write a block of as many as 1300 words, which could produce (if all words were numeric) 13,000 tape characters.

By overlaying location ABLOCKMAX with a word of the form

+000000XXXX

the user will cause Sort 90 to create tape blocks of not more than XXXX words. Naturally this figure must not be smaller than the number of words making up the largest possible data record during the sort.

SORT 90 MODIFICATION I-6

Main Program

Purposes:

Editing of Phase I input records; selective deletion of some Phase I input records; any other record-by-record processing in Phase I.

Label of Exit
Instruction

Labels of Return
Instruction

ARECDBRNCH

AWRITE
ADELETE

Description:

Instruction ARECDBRNCH is executed once for every input record in Phase I, immediately before the record is to enter the Phase I internal sorting routine. It may be overlaid with a Branch to added programming. Whenever the program passes through ARECDBRNCH, digits 2-5 of index word ARECORD contain the start address of the record. ARECORD may thus be used to index references to fields within the record, or to control movements of the record. Its contents may not be disturbed.

The current input record may be edited, according to the restrictions described on page 52. The added programming should return to AWRITE when the editing of an input record has concluded. All control data field descriptions on the Sort 90 Second Control Card should reflect the positions of these fields within the edited record, rather than the input record. If editing actions alter either the record type or the record length, or both, columns 66-69 of the Sort 90 Control Card 1 must be punched.

If after examination of an input record it is desired to delete it from the sort, the added programming should return to ADELETE.

Any other record-by-record processing may be done by exiting from ARECDBRNCH to added programming, and the return instruction except for deletions is always AWRITE.

The Phase I hash total, if any, is taken after re-entrance is made to AWRITE or ADELETE; if Phase I editing is performed the Sort 90 First Control Card hash field description should describe the position of the field in the edited record. If hash total option 1, 4, 5, or 7 is used, and if the editing of records involves changing the hash total field position, and if some records are deleted — if all three conditions apply, the hash total fields of the records to be deleted should be re-positioned before return is made to ADELETE. If Phase I editing involves changing the contents of the hash total field, hash total options 1, 4, 5, and 7 may not be used.

SORT 90 MODIFICATION I-7

Main Program

Purpose:

To alter the IOCS treatment of the sort input file through modification of its DTF entry.

Description:

The first word of the nine-word sort input file DTF table occupies location IOCSFTBL01. Changes may be made to the DTF fields listed below. Since these DTF fields are partial-word fields, the user must locate the Phase I

load card containing the word altered, and transfer to a patch card word all digits which are to remain unaltered in the DTF.

The DTF fields which may be altered are:

<u>Line</u>	<u>Label (for reference)</u>	<u>Present Contents</u>
12	OPENPROC	2
13	CLSEPROC	3
30	RWDPROC	3

All other fields must not be changed, either because their present contents are required by the logic of Sort 90 or because its Phase I assignment routine extracts the pertinent information from Sort 90 control cards and inserts it into the DTF fields.

The EOFPROCD exit is used by the sort, but is followed by an exit instruction; see Modification I-9.

The EORPROCD exit is used by the sort; added end-of-reel coding may however be entered by overlaying location AENDOFREEL with a branch to the desired routine; return to the sort is made with the instruction

B 0+IOCSIXF

The instruction in AENDOFREEL is used as a Branch-NOP switch, and consequently any overlaid instruction must be a Branch (operation code +01).

SORT 90 MODIFICATION I-8

Main Program

Purpose:

To modify the IOCS treatment of the sort input file tape labels.

Description:

The first word of the eight-word input file DC label information entry occupies location AINLABEL. Changes may be made as detailed below.

Word AINLABEL itself contains the input file mask. Digits 0 - 4 may be modified as desired; digits 5 - 9 must remain "00011."

Word AINLABEL+1 contains a Branch address to the Sort "Exit 6" routine. The routine checks (if record count option 1 or 3 is elected) the computed record count for this reel against the record count in the input trailer label, and checks (if hash total option 1, 4, 5, or 7 is elected) the computed hash total for this reel against the hash total in the input trailer label. If no checks are made, if the checks indicate no discrepancies, or if a discrepancy has occurred but the operator elects either to ignore it or to delete the entire reel from the sort, the program will immediately pass through instruction AMORETRLR, which may be overlaid with a Branch to added coding to process the trailer label, the first word of which occupies location IOCSLBAREA+1. Return should be made to location AMORETRLR+1.

If record count and/or hash total discrepancies cause the operator to reprocess the reel, AMORETRLR will not be active for this occurrence of trailer label reading for the reel.

Word AINLABEL+2 (Exit 7) contains a Branch address to an instruction in the sort which merely returns control to IOCS. A different branch address may be inserted in this word for additional header label processing; return should be made to IOCS by the instruction

B 0+IOCSIXF

The alphabetic fields contained in AINLABEL+3 through AINLABEL+7 are filled from fields contained on the Input Label Control Card, if the latter fields are non-zero. The user has two methods of establishing the fields against which the contents of input file header labels will be checked: entering the fields through the Input Label Control Card, or overlaying the fields directly in the DC entry. If the input file is labeled, however, an Input Label Control Card must be furnished. Control over what header label checks are to be made is gained through modification of the mask in AINLABEL.

SORT 90 MODIFICATION I-9

Main Program

Purpose:

To cause the reading or writing of one or two additional tape files during Phase I.

Description:

To read or write an additional tape file in Phase I, the user must assemble a main program modification consisting of a DTF entry, a label DC information entry, a DA for record storage areas, the requisite IOCS macro-instructions (including OPEN and CLOSE), the required processing instructions, and the required linkages from and into the sort program. Exit instruction ARECDBRNCH (see Modification I-6) will be involved for the purposes of examining and processing records of the input file which have relation to the records of the additional file(s) being read or written. If one additional file is used, its DTF table should occupy the nine words beginning with location IOCSFTBL03; if two additional files are used, their DTF tables should occupy the two nine-word areas beginning with IOCSFTBL03 and IOCSFTBL04, respectively.

The information sheet accompanying the Sort 90 program deck will give the actual addresses of all IOCS labels required during the Autocoder assembly of the modification. These are IOCSIRTAIN, IOCSIGEN, IOCSIOPEN, IOCSICLOSE, IOCSIXF, and IOCSIXG. It is not possible to use IOCS checkpoint and restart routines for additional Phase I files.

The priority digits (PRIORITY field of the DTF) used for the additional files must be selected from the following: 0, 3, 4, 5, 6, 7, 8, and 9.

The OPEN instruction for any Phase I file must be issued through exit AADASSIGN (see Modification I-3), and the CLOSE instruction through exit AWINDUP (see Modification I-10).

Purpose:

End-of-phase processing in Phase I.

**Label of Exit
Instruction**

Placement and Use

AWINDUP Instruction AWINDUP occurs in the Phase I end-of-phase routine after the last Phase I output records have been written, and before Phase I output tapes are CLOSED. An added routine originating in AWINDUP should return to instruction AWINDUP+1. This exit should be used for performing the final processing of and CLOSEing additional tape files, if any, read or written during Phase I.

AENDPHASE1 Instruction AENDPHASE1 occurs at the beginning of a routine to make linkage to the load program for purposes of loading Phase II. It and the next three storage locations contain the following:

```

AENDPHASE1      B      *+1
                  PC      *+2
                  B      0000
                  DC
                  +1111111111
    
```

If a loading system is used which does not leave a "pivot" to itself in location 0000, AENDPHASE1+1 should be used as an exit to the loading procedure.

Note: If after interrupting Phase II and later restarting, Phase I discovers that it has been reloaded, Phase III being desired, Sort 90 will branch directly to AENDPHASE1+1.

PHASE II - MODIFICATIONS

Purpose:

Miscellaneous assignment routine modifications in Phase II.

**Label of Exit
Instruction**

Placement and Use

BASSIGN This is the first instruction of the Phase II assignment routine. Any general assignment modification may be entered by overlaying BASSIGN by a Branch to the modification. The sort program should be re-entered at BASSIGN+1. Exit BASSIGN is available only before the first Phase II pass, as the assignment routine is destroyed before the first pass begins.

Purpose:

Exits for beginning-of-pass and end-of-phase processing in Phase II.

Symbolic Label

Placement and Use

BSTARTPASS This instruction is the first instruction of the Phase II initialization routine, and is executed at the beginning of each Phase II pass. Any general housekeeping modification may be entered through this instruction; return should be made to location BSTARTPASS+1.

BENDPHASE2 When Phase II concludes its last pass, it enters a routine to make linkage to the load program to load Phase III. Instruction BENDPHASE2 may be overlaid with a Branch to added coding such as a specialized load initiator. The contents of BENDPHASE2 and the next three locations are:

BENDPHASE2	B	*+1
	PC	*+2
	B	0
	DC	+111111111

Note: If, following a Phase II interruption and later restart, Phase II discovers that it has been loaded, Phase III being desired, it will branch to BENDPHASE2+1. The link to the loading procedure must therefore begin with this instruction.

Purpose:

Processing of records on a serial basis in Phase II passes.

Label of Exit
Instruction

Label of Return
Instruction

BRECEXITBC

BPUT

Description

This exit provides for any simple record-by-record processing in Phase II passes, such as the accumulation of various control totals. It is not active in a pass in which Phase II summarizing is being performed. The instruction contained in BRECEXITBC is located elsewhere when executed.

When using the exit, index word BRECORD contains in digits 2-5 the start address of the "current" record being processed. Return must always be to BPUT.

It is desired to execute different processing routines during various Phase II passes, the initialization routine exit BSTARTPASS (see Modification II-4) may be used to store into BRECEXITBC a Branch instruction appropriate to the routine for the particular Phase II pass. If BRECEXITBC has previously been activated and it is desired to de-activate it, simply store the instruction

B BPUT

SORT 90 MODIFICATION II-4

Main Program

Purpose:

Summarization in Phase II of records with identical control data.

Label of Exit
Instruction

Labels of Return
Instructions

BSUMBRANCH

BSUMMARIZE
BWRITE

The instruction BSUMBRANCH is traversed whenever two records with identical control data are discovered in Phase II, and if the Phase II summarizing pass control (see Sort 90 Second Control Card columns 66-67 and Modification II-5) indicates that summarizing ability applies to the Phase II pass currently being executed.

When the program reaches BSUMBRANCH, digits 2-5 of index word BBASERECRD contain the start address of one record, and digits 2-5 of index word BRECORD contain the start address of a second record the control data of which is identical to that of the first record.

By means of instructions indexed by these index words, data fields (such as amount fields) of the BRECORD record may be added to the corresponding fields of the BBASERECRD record. The routine to perform this addition should be entered by a Branch instruction which overlays BSUMBRANCH, and should when complete re-enter the sort program at BSUMMARIZE. The BBASERECRD record will be retained; the current BRECORD will be deleted from the sort, and a new BRECORD record will be obtained as above.

The summarizing instructions may not affect any of the following fields in any record: all control data fields, the hash total field (if one is specified), and the length field of Form 3 records.

SORT 90 MODIFICATION II-5

Main Program

Purpose:

Pass-by-pass control over Phase II summarizing.

Label of Exit
Instruction

Labels of Return
Instructions

BSUMMTEST

BDOSUMM
BSKIPSUMM

In the initialization for each Phase II pass, a comparison is made between the "summarizing pass control" as punched in the Sort 90 Second Control Card and the number of Phase II passes which remain to be executed. When the first quantity exceeds or equals the second, the current Phase II pass is established with the ability to detect summarizable records. When the second quantity exceeds the first, the current Phase II pass does not detect summarizable records.

This procedure may be overridden through use of a decision making routine which is entered by overlaying BSUMMTEST with a Branch to the routine. If the sort is re-entered at BDOSUMM, the current Phase II pass will detect summarizable records; if the sort is entered at BSKIPSUMM, it will not.

When BSUMMTEST is executed, Accumulator 2 contains in its low-order positions the number of expected further Phase II passes (including the pass just to be entered). This quantity may be examined as desired; Accumulator 2 may not however be disturbed.

SORT 90 MODIFICATION II-6

Main Program

Purpose:

To cause index word, electronic switch, and main storage areas used by Phase II modifications to be written in the IOCS checkpoint record for proper interruption and later restart of Phase II.

Description:

A list of eight RDWs, as follows, is contained in Phase II for the use of IOCS in writing a checkpoint record:

BCHPTRDWS	DRDW	(A plus RDW defining the Phase II index word area.)
	DRDW	(A plus RDW defining the Phase II main storage area. The stop address is taken from columns 48-51 of Control Card 1.)
	DRDW	+, *
	DRDW	+, *
	DRDW	+, *
	DRDW	+, *
	DRDW	+, *
	DRDW	-, *

The first two RDWs may not be disturbed, but the last six are "dummy" RDWs and should be overlaid as necessary with RDWs descriptive of storage areas used by Phase II modifications. All of the RDWs should be signed plus, except for the last, which should be signed minus.

PHASE III - MODIFICATIONS

SORT 90 MODIFICATION III-1

Assignment Routine

Purpose:

Miscellaneous exit points during Phase III assignment.

Description:

Following are two instructions in the Phase III assignment routine from which added programming may be entered.

Label of Exit
Instruction

Placement and Use

CASSIGN	This is the first instruction of the Phase III assignment routine. Any general housekeeping routines may be entered; return should be made to CASSIGN+1.
CADDASSIGN	This is the last instruction of the Phase III assignment routine; again, any general housekeeping modifications may be entered from it. Return should be made to CGENERATE, following which the generation of RDWs and read-in of records will destroy the Phase III assignment routine. Since the IOCS OPEN subroutine is also destroyed, OPENing of additional tape files in Phase III should be accomplished through this exit.

SORT 90 MODIFICATION III-2

Main Program

Purposes:

Editing, deleting, and summarizing of records in Phase III; insertion of new records into the sort output file; other modifications involving processing of records on a one-by-one basis in Phase III.

Labels of Exit
Instructions

Labels of Return
Instructions

CRECEXIT	CSEQCHECK
CHIBRANCH	CWRITE
CSUMBRANCH	CDELETE
CSTEPDOWN	CSUMMARIZE
	CINSERT (subroutine)

Description:

These modifications center around the sequence-check made in Phase III. The sequence-check is a comparison of the control data of the current "low" record, the start address of which occupies digits 2-5 of index word CRECORD, with the control data of the last sort output record, the start address of which occupies digits 2-5 of index word CBASERECRD. The flow-

chart (Figure 5) illustrates the positionings of the various exit and return points; CSEQCHECK is the label of the first instruction of the sequence-check routine itself, and may be used as a return instruction only if exit has been made from CRECEXIT. Control passes through CSUMBRANCH when the contents of the number of control fields indicated in column 65 of Sort 90 Control Card 2 are equal; see page 32.

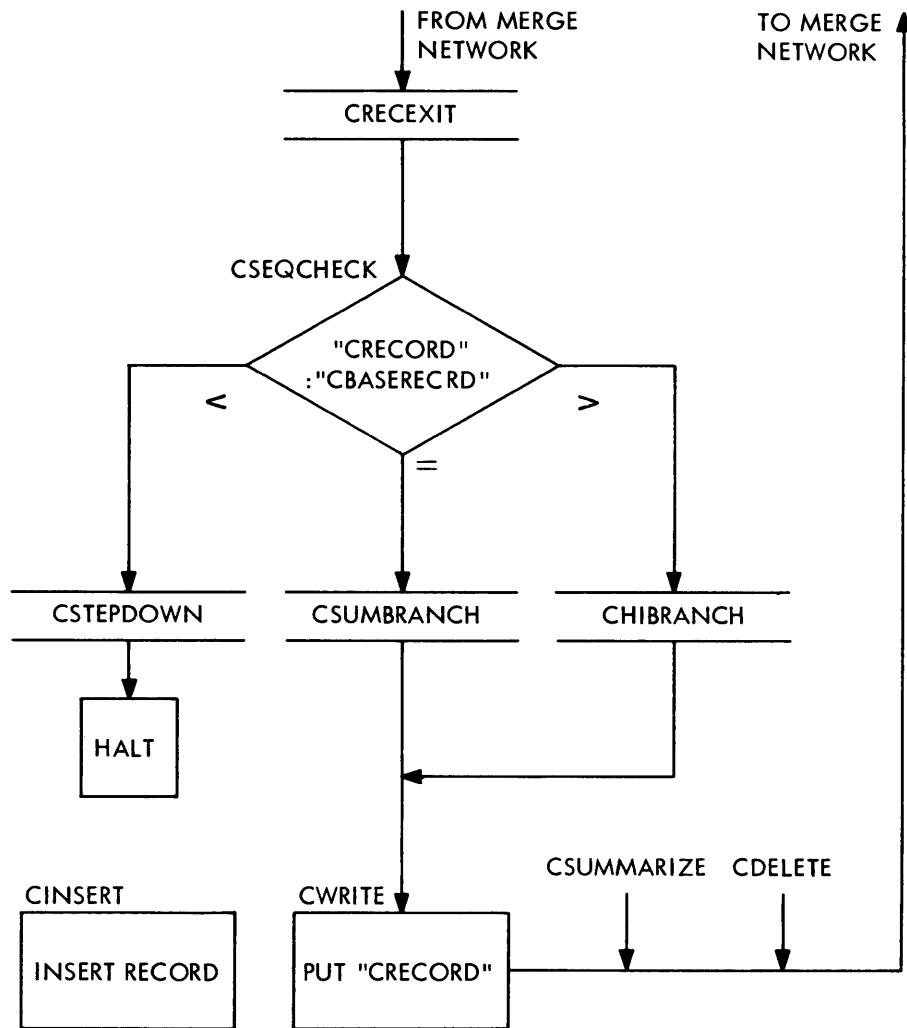


Figure 5

At the beginning of Phase III, there will be no CBASERECRD record until the CWRITE routine (or the CINSERT subroutine) has been executed once; a switch within the CSEQCHECK routine sends control through CHIBRANCH until this time. The very last record from the Phase III input tapes is available through any of the above exits as a CRECORD record, but not as a CBASERECRD record; see Modification III-6.

Return instructions are employed according to the actions performed or to be performed upon the CRECORD record.

Deleting. If from any of the exits return is made to CDELETE, the current CRECORD record will be dropped from the sort. The current CBASERECRD record and the next CRECORD will be involved in the next sequence-check.

Summarizing. Summarizing routines should originate from exit CSUMBRANCH; data fields of the CRECORD record may there be added to those of the CBASERECRD record, and return made to CSUMMARIZE. This return will cause the current CRECORD to be dropped and the next CRECORD record to be compared with the current CBASERECRD, with control going through CSUMBRANCH if its control data is again identical. It can be seen that if there is a string of equal records, the first will result in control going through CHIBRANCH, and all others through CSUMBRANCH.

Inserting. At any point in added coding, use or multiple use may be made of the CINSERT subroutine, which will move to the sort output file a designated record, insert its start address in digits 2-5 of CBASERECRD, and return control to the added coding. The calling sequence is:

```
BLX          CINSERTX, CINSERT
DRDW         +xxxx, yyyy
```

where the RDW defines the start-stop addresses of the record to be inserted; the CINSERT subroutine will return to the instruction which follows the DRDW entry. The record which was the CBASERECRD prior to use of the CINSERT subroutine is made inaccessible through use of the subroutine, and may not be referenced by further instructions in the added coding. If the contents of the control data fields of the inserted record(s) bears no resemblance to those of further sort output records, and if the routine in which the records are inserted does not return control to CWRITE, the sequence-check routines must not be permitted to compare the next CRECORD record with the last of the inserted records. If the inserted records are inserted in sequence with the records of the sort output file (for example, if a new file were being merged with the Phase III output file), the sequence-check routine should not be bypassed as in this case it has the additional function of verifying the correctness of the insertions.

Editing; other processing. For editing and other processing which does not involve the deletion of the CRECORD record, return should be made to CWRITE. (This is the normal return; CSUMBRANCH and CHIBRANCH transfer to CWRITE unless modified.)

The CRECORD record will be made the CBASERECRD record, and a new CRECORD record obtained for the next cycle. Editing of the CRECORD record must follow the restrictions mentioned on page 52. In addition, if editing involves changes in the positions or contents of control data fields, the entire sequence-check routine must be replaced. This may be done by:

1. Exiting from CRECEXIT
2. Editing the CRECORD record.
3. Comparing control data of the (just-edited) CRECORD record with that of the (previously-edited) CBASERECRD. (A first-time-only switch should be used to bypass this comparison for the first cycle.)
4. Returning to CWRITE.

If the position or contents of the hash total field (if any) is altered by Phase III editing, hash total options 2, 4, 6, and 7 may not be used. If hash total option 3 is used, the hash total field position indicator in the Sort 90 First Control Card must refer to the record as edited in Phase III.

SORT 90 MODIFICATION III-3

Main Program

Purpose:

To alter the IOCS treatment of the sort output file through modification of its DTF entry.

Description:

The first word of the nine-word sort output file DTF table occupies location IOCSFTBL01. Changes may be made to the two DTF fields listed below. Since both are partial-word fields, the user must locate the Phase III load card containing the word affected, and transfer to a patch card all digits of the word which are to be unaltered.

The two DTF fields which may be altered are:

<u>Line</u>	<u>Reference Label</u>	<u>Present Contents</u>
13	CLSEPROC	3
30	RWDPROCD	3

All other fields may not be changed, either because their present contents are required by the logic of Sort 90 or because the Phase III assignment routine inserts into them the proper information. The EORPROCD exit is used by Phase III, but an exit is provided at the termination of the sort FOR routine for added coding. To use this exit, overlay instruction CEOREXIT with a Branch to the added routine, and re-enter IOCS by the instruction

B 0+IOCSIXF

It is possible to include as Phase III modifications IOCS routines which are not contained within Sort 90. A convenient procedure for such modifications is the assembly of a program consisting of a DTF entry which occupies the same locations as the sort output file DTF entry and the desired IOCS macro-instructions or routines; after assembly, the cards containing the DTF entry and the generated file scheduler are destroyed. Relevant locations pertaining to the sort output file are: IOCSFTBL01, first word of file scheduler; COUTPUTXWA, the first output file index word; and COUTPUTXWB, the second sort output file index word. Knowledge of the addresses of other IOCS symbolic locations which are loaded with Sort 90 is also required for proper assembly. This information is included in the material which accompanies the Sort 90 program deck.

Purpose:

To modify the IOCS treatment of the sort output file tape labels through modification of its DC label information entry.

Description:

The first word of the twenty-one word output file DC label information entry occupies location COUTLABEL. Changes may be made as detailed below.

Word COUTLABEL itself contains the output file mask. Digits 0-4 may be modified as desired; digits 5-9 must remain 11111.

Word COUTLABEL+1 contains a Branch address to the Phase III "Exit 1" routine and should not be altered. The routine inserts the record count into the trailer label, the first word of which occupies location IOCSLBAREA+1 and inserts into the trailer label the hash total, both these actions being conditioned by the options chosen as to record counts and hash totals. The program then passes through location CENDEXIT1, which may be replaced with a Branch to added coding; return should then be made to IOCS with the instruction

B 0+IOCSIXF

Words COUTLABEL+2, COUTLABEL+3, COUTLABEL+4, and COUTLABEL+5 ("Exit 2," "Exit 3," "Exit 4," and "Exit 5" respectively) all contain addresses of an instruction within Phase III which merely sends control back to IOCS. One or more of these words may be replaced with a Branch address to added routines; in all cases added routines should return control to IOCS by the instruction

B 0+IOCSIXF

Words COUTLABEL+6 through COUTLABEL+20, the alphabetic portion of the output file DC entry, will be filled (in part) by Phase III from non-zero fields on the Output Label Control Card, and may also be overlaid or established by added routines. Control over what checks are to be made and what fields are to be regulated by IOCS is gained through modification of the mask in COUTLABEL.

Purpose:

To cause the reading or writing of additional tape files during Phase III.

Description:

To read or write an additional file in Phase III, the user must assemble a main program modification consisting of a DTF entry, a label DC information entry, a DA for record storage areas, the requisite IOCS macro-

instructions (including OPEN and CLOSE), and the required linkages to and from the sort program. These linkages will normally involve exit and return instructions as detailed in Modification III-2.

The number of additional Phase III files which may be used is equal to (7-M), where M is the order of merge: With a three-way merge, four files could thus be processed. Following is a list of areas available for DTFs for the added files, with an indication of which priority digits (PRIORITY field of the DTF) may not be used.

<u>Label of First Word of DTF</u>	<u>Priority Digit</u>	<u>Use</u>
IOCSFTBL01	1	Sort 90 output file.
IOCSFTBL02	3	First Phase III input file.
IOCSFTBL03	4	Second Phase III input file.
IOCSFTBL04	5	Third Phase III input file; may be overlaid with a new DTF if order of merge is 2.
IOCSFTBL05	6	Fourth Phase III input file; may be overlaid with a new DTF if order of merge is 2 or 3.
IOCSFTBL06	7	Fifth Phase III input file; may be overlaid with a new DTF if order of merge is 2, 3, or 4.
IOCSFTBL07		Vacant; may always be used for new DTF.
IOCSFTBL08		Vacant; may always be used for new DTF.

Added DTF entries may be placed in any of the unused tables listed above, but it is suggested that those at the end of the list be used first.

The OPEN macro-instruction for added files must be given by a routine originating at exit CADDASSIGN (see Modification III-1); exit CWINDUP (see Modification III-6) is suggested as an exit point for processing of the terminal records of added files and for issuance of a CLOSE macro-instruction referencing the files.

The information sheet accompanying the Sort 90 program deck will give actual addresses of all IOCS labels required during Autocoder assembly of the modification. These are IOCSIRTAIn, IOCSIGEN, IOCSIOPEN, IOCSICLOSE, IOCSIXF, and IOCSIXG. Checkpoint and restart routines used by Phase III will automatically apply to additional Phase III files.

SORT 90 MODIFICATION III-6

Main Program

Purpose:

Exits for end-of-phase processing in Phase III.

Description:

Following are two instructions in the Phase III end-of-job routine from which added coding may be entered:

Instruction

Placement and Use

CWINDUP This instruction occurs after the last sort output record has been PUT, and before the last Phase III output tape has been CLOSED. It may therefore be used to enter any general end-of-job routines; return should be made to CWINDUP+1. A particular use of exit CWINDUP, if Modification III-5 is activated, is the completion of processing and the closing of additional Phase III input or output files. When control goes through CWINDUP, the last Phase III record is available as the CBASERECRD record, and the CINSERT subroutine may be used (see Modification III-2).

CENDPHASE3 When Phase III is completely finished, it comes to a programmed halt (Halt 3333); if START is depressed, transfer is made to location IOCSEJLOOP, in the Input/Output Control System, to set up a "busy loop" for SPOOL:

```
CENDPHASE3  B    *+1
              HP   3333(0)
              B    IOCSEJLOOP
```

Instruction CENDPHASE3 may be overlaid with a Branch to added coding before the halt, such as a routine for loading the next program. The IOCSEJLOOP routine assumes a 10,000 word 7070; if a 5,000 word 7070 is used, alter digits 6-9 of IOCSEJLOOP, IOCSEJLOOP+5, and IOCSEJLOOP+6 to "4999."

SORT 90 MODIFICATION III-7

Main Program

Purpose:

To cause index word, electronic switch, and main storage areas used by Phase III modifications to be written in the IOCS checkpoint record for proper interruption and later restart of Phase III.

Description:

A list of eight RDWs, as follows, is contained in Phase III for the use of IOCS in writing a checkpoint record:

CCHPTRDWS	DRDW	(A plus RDW defining the Phase III index word area.)
	DRDW	(A plus RDW defining the Phase III main storage areas. The stop address is taken from columns 52-55 of Control Card 1.)
	DRDW	+*, *
	DRDW	+*, *

DRDW	+*, *
DRDW	+*, *
DRDW	+*, *
DRDW	-, *

The first two RDWs may not be disturbed, but the last six are "dummy" RDWs and should be overlaid as necessary with RDWs descriptive of storage areas used by Phase III modifications. All of the RDWs should be signed plus, except for the last, which should be signed minus.

APPENDIX I

GLOSSARY OF TAPE SORTING AND TAPE HANDLING TERMS

BLOCKING

The incorporation of several data records, or item records, into one long tape record, or block, in order to reduce tape space and tape acceleration time per data record. The number of data records per block is the blocking factor of a file.

CHECKPOINT

The writing of the contents of 7070 storage on tape, to permit later restart.

COLLATING SEQUENCE

An ordering of characters or digits based upon their effects in compare operations.

CONTROL DATA

That portion of a record which is compared with the corresponding portion of another record for sorting purposes. Control data consists of one or more control data fields. That part of a control data field which lies within boundaries of a 7070 word is termed control data segment.

HASH TOTAL

A total taken on an arbitrary field of records being sorted, for control purposes only. The hash total taken during one pass may be compared with that taken during a previous pass to aid in detecting any mishandling of records.

MERGE

A sorting procedure consisting of the dovetailing together of records from two or more sequences in order to produce one sequence. A merge executed totally within a computer is called an internal merge. The number of sequences being merged together is called the order of merge. A merge of order M is also called an M-way merge.

PASS

One complete cycle of reading, processing, and writing an entire file.

RECORD COUNT

A simple count of the records processed during one pass; used for control purposes.

RESTART

The resumption of a program following a previous interruption. A restart is usually effected by reading back into storage a previously written check-point record.

SEQUENCE

A string of records arranged in the desired sorting order; they are in order by control data according to the collating sequence being used.

STEPDOWN

A sequence break, or condition obtaining between two records the second of which is not in sequence with the first.

APPENDIX II

CALCULATION OF SORTING TIMES

The procedure for calculating the running time of Sort 90 for any one sorting application is similar to that for calculating the running time for any computer program with magnetic tape input and output. The tape input/output time and the process times are computed, and the larger of these times, incremented by any secondary factors such as rewind time, is taken as the execution time of the program. In the case of Sort 90, which is a multi-pass program, this procedure must be performed for each pass and the resultant times added.

The step-by-step procedure described has the important advantage of yielding a very precise time estimate, in spite of the large number of variable and occasionally unknown factors which may contribute to the running time of Sort 90. It furthermore allows unusual or non-standard conditions to be taken into account, e. g., the effect upon process time of additional programming which modifies or extends Sort 90.

Sort 90 Timing Calculations

The step-by-step procedure for calculating Sort 90 running times makes use of eight tables. These tables are presented below, along with a discussion of their proper use. A description of the procedure itself follows on page 82.

The time figures which are contained in the tables associated with the timing procedure refer to a 7070. The step-by-step procedure allows the timing figures to be adjusted for the operation of Sort 90 on a 7074.

In some cases a value to be used in looking up an entry in a table will not be listed in the table. In this event, the closest table value should be used, or, preferably, an interpolation with the surrounding values should be made and used to extract the information from the table.

Tables I and II differ only in that Table I is to be used with a 5,000-word 7070 or 7074, and Table II with a 10,000-word 7070 or 7074. Each of these tables gives Phase I basic process times, in seconds per 1,000 records, tabulated against the length, in words, of a data record being sorted. Table III gives basic process times for one Phase II pass and for Phase III, tabulated against the order of merge; these times are independent of data record length and machine size.

Adjacent to each basic process time in Tables I, II, and III is a "time for each additional comparison." The basic process times have been calculated on the assumption that one comparison will rank any two data records of the file being sorted, i. e., that the major control data segments are unequal. The configuration of the records being sorted may be such that normally two records cannot be ranked until more than one comparison has been performed. The process time required for additional comparisons is obtained by multiplying the number of needed additional comparisons (see "Optimal Control Data Arrangement," page 16) by the time for each additional comparison.

TABLE I
 SORT 90 PHASE I PROCESS TIMES IN SECONDS PER 1,000 RECORDS
 (5,000-word 7070)

Record Length (words)	Phase I Basic Process Time	Phase I Time for Each Additional Comparison	Phase I Output Sequence Length	Sort Blocking Factor			
				2-way merge	3-way merge	4-way merge	5-way merge
1	5.72	2.25	333	208	156	125	104
2	5.15	2.00	238	138	104	83	69
3	5.15	2.00	185	104	78	62	52
4	4.58	1.75	151	83	62	50	41
5	4.58	1.75	128	69	52	41	34
6	4.58	1.75	111	59	44	35	29
7	4.58	1.75	98	52	39	31	26
8	4.58	1.75	87	46	34	27	23
9	4.00	1.50	79	41	31	25	20
10	4.00	1.50	72	37	28	22	18
12	4.00	1.50	61	32	24	19	16
14	4.00	1.50	53	27	20	16	13
16	4.00	1.50	47	24	18	14	12
18	4.00	1.50	42	21	16	13	10
20	3.41	1.25	38	19	14	11	9
24	3.41	1.25	32	16	12	10	8
28	3.41	1.25	28	14	10	8	7
32	3.41	1.25	24	12	9	7	6
36	3.41	1.25	22	11	8	6	5
40	2.82	1.00	20	10	7	6	5
50	2.82	1.00	16	8	6	4	4
60	2.82	1.00	13	6	5	4	3
70	2.82	1.00	11	5	4	3	2
80	2.19	0.75	10	5	3	3	2
90	2.19	0.75	9	4	3	2	2
100	2.19	0.75	8	4	3	2	2
150	1.52	0.50	5	2	2	1	1
200	1.52	0.50	4	2	1	1	1
250	1.52	0.50	3	1	1		
300	0.73	0.25	2	1	1		
400	0.73	0.25	2	1			

TABLE II
 SORT 90 PHASE I PROCESS TIMES IN SECONDS PER 1,000 RECORDS
 (10,000-word 7070)

Record Length (words)	Phase I Basic Process Time	Phase I Time for Each Additional Comparison	Phase I Output Sequence Length	Sort Blocking Factor			
				2-way merge	3-way merge	4-way merge	5-way merge
1	6.29	2.50	998	624	468	374	312
2	6.29	2.50	713	416	312	249	208
3	5.72	2.25	554	312	234	187	156
4	5.72	2.25	453	249	187	149	124
5	5.72	2.25	384	208	156	124	104
6	5.72	2.25	332	178	133	107	89
7	5.15	2.00	293	156	117	93	78
8	5.15	2.00	262	138	104	83	69
9	5.15	2.00	237	124	93	74	62
10	5.15	2.00	217	113	85	68	56
12	5.15	2.00	184	96	72	57	48
14	5.15	2.00	161	83	62	49	41
16	4.58	1.75	142	73	55	44	36
18	4.58	1.75	128	65	49	39	32
20	4.58	1.75	116	59	44	35	29
24	4.58	1.75	97	49	37	29	24
28	4.58	1.75	84	43	32	25	21
32	4.00	1.50	74	37	28	22	18
36	4.00	1.50	66	33	25	20	16
40	4.00	1.50	60	30	22	18	15
50	4.00	1.50	48	24	18	14	12
60	3.41	1.25	40	20	15	12	10
70	3.41	1.25	34	17	13	10	8
80	3.41	1.25	30	15	11	9	7
90	3.41	1.25	27	13	10	8	6
100	3.41	1.25	24	12	9	7	6
150	2.82	1.00	16	8	6	4	4
200	2.82	1.00	12	6	4	3	3
250	2.19	0.75	9	4	3	2	2
300	2.19	0.75	8	4	3	2	2
400	2.19	0.75	6	3	2	1	1
500	1.52	0.50	4	2	1	1	1

TABLE III
SORT 90 PHASE II AND PHASE III PROCESS TIMES IN SECONDS PER 1,000 RECORDS

	Order of Merge			
	2-way	3-way	4-way	5-way
Phase II Basic Process Time (One Pass)	1.12	1.28	1.44	1.60
Phase II Time for each Additional Comparison (One Pass)	0.25	0.40	0.50	0.63
Phase III Basic Process Time	1.46	1.75	2.03	2.32
Phase III Time for each Additional Comparison	0.50	0.63	0.75	0.88

Note that all control data segments are not usually involved in a comparison between two data records; only a small number of control data segments are involved. (The determination of which control data segments are involved requires some knowledge of the actual control data contained in the records.)

If any phase of the Sort is modified to perform additional functions such as editing or summarizing, the execution time of the added programming should be determined, as it contributes to the total process time of that phase.

Tables I and II give the number of records in a Phase I sorted output sequence, based upon the number of words in a record. The total number of sorted sequences which Phase I will produce is equal to the total number of records being sorted divided by the number of records in one sorted sequence. Tables I and II also give the blocking factor (number of data records in a tape block, or tape record) used by the sort throughout its operation; this figure is based upon both the record length in words and the order of merge being used.

Table IV shows the number of Phase II passes necessary in a given run of Sort 90, tabulated against the number of sorted sequences produced by Phase I and the order of merge being used.

Tables V, VI, VII, and VIII give the time, in seconds per 1,000 records, required to read or write data records on magnetic tape. The four tables show input/output times based upon tape unit model, i. e., 729 II and 729 IV, and tape density, i. e., high (556 characters per inch), and low (200 characters per inch). The input/output time is tabulated against the blocking factor and the "average" number of tape characters required to express a data record on tape. The "average" number depends on the average number of words in a data record (if Form 2 or Form 3); the number of alphabetic and number of numerical words; the number of mode change characters; and, if zero suppression is used, the number of high-order zeros in numerical words.

Tables V, VI, VII, and VIII are used to determine three different tape times; the Phase I input tape time, the sort tape time (Phase I output, Phase II input and output, Phase III input), and the Phase III output tape time.

The Phase I input and Phase III output tapes may be written with either high or low density, and with or without zero suppression. On the other hand, all sort work tapes are written with high density and zero suppression.

It should be noted in connection with tape rewind times that full rewinds are not usually performed on sort work tapes. The total number of records is spread over a number of tape reels equal to the order of merge, so that any one tape reel will not be full.

The step-by-step procedure for calculating Sort 90 running times is presented below. The result of performing each step should be entered in the space provided to the right of each step; subsequent steps will refer to the results of previous steps by the step number.

TABLE IV
REQUIRED NUMBER OF PHASE II PASSES

Number of Phase I Output Sequences	Order of Merge			
	2-way	3-way	4-way	5-way
1-2	0	0	0	0
3	1	0	0	0
4	1	1	0	0
5	2	1	1	0
6-8	2	1	1	1
9	3	1	1	1
10-16	3	2	1	1
17-25	4	2	2	1
26-27	4	2	2	2
28-32	4	3	2	2
33-64	5	3	2	2
65-81	6	3	3	2
82-125	6	4	3	2
126-128	6	4	3	3
129-243	7	4	3	3
244-256	7	5	3	3
257-512	8	5	4	3
513-625	9	5	4	3
626-729	9	5	4	4
730-1024	9	6	4	4
1025-2048	10	6	5	4
2049-2187	11	6	5	4
2188-3125	11	7	5	4
3126-4096	11	7	5	5
4097-6561	12	7	6	5
6562-8192	12	8	6	5
8193-15625	13	8	6	5
15626-16384	13	8	6	6
16385-19683	14	8	7	6
19684-32768	14	9	7	6
32769-59049	15	9	7	6
59050-65536	15	10	7	6
65537-78125	16	10	8	6

STEP-BY-STEP PROCEDURE

1. Enter the following information:
 - a. For Form 1 records, data record length in words; Form 2, maximum data record length in words; Form 3, average data record length in words. _____
 - b. Average number of tape characters per data record. _____
 - c. For Form 1 and 2 records, the input file blocking factor; for Form 3 records, the maximum number of words in an input tape block divided by (1a). _____
 - d. For Form 1 and 2 records, the desired output file blocking factor; for Form 3, the desired maximum number of words in an output tape block divided by (1a). _____
 - e. Number of data records in the file. _____
 - f. Number of needed additional comparisons. _____
2. Using (1a), extract from Table I or II:
 - a. Phase I basic process time. _____
 - b. Phase I time for each additional comparison. _____
 - c. Phase I output sequence length. _____
 - d. Sort blocking factor. _____
3. Multiply (1f) by (2b). _____
4. Enter the time for additional Phase I programming. _____
5. If Form 3 records are being sorted, enter 0.52. _____
6. If hash total option 1, 2, 4, 5, 6, or 7 is specified, enter 0.21. _____
7. Enter the Phase I process time: for a 7070, the sum of items (2a), (3), (5), and (6); for a 7074, the sum of these items divided by 5. _____
8. Using (1b) and (1c), extract from Table V, VI, VII, or VIII the Phase I input tape time. _____
9. Using (1b) and (2d), extract from Table V, or VII the sort tape time. _____
10. Enter the total Phase I time: If two tape channels are used, enter the largest of the three items (7), (8), and (9); if one tape channel is used, add (8) and (9) and enter their sum or (7), whichever is larger. _____

11. Extract from Table III:
 - a. Phase II basic process time for one pass. _____
 - b. Phase II time for each additional comparison. _____
12. Multiply (1f) by (11b). _____
13. Enter the time for additional Phase II programming. _____
14. If Form 3 records are being sorted, multiply (1a) by 0.024, add 0.19 to their product, and enter the result. _____
15. If hash total option 2, 4, 6, or 7 is specified, enter 0.21. _____
16. Enter the Phase II process time for one pass: for a 7070, the sum of items (11a), (12), (13), (14), and (15); for a 7074, the sum of these items divided by 5. _____
17. Enter the Phase II time for one pass: if two tape channels are used, enter the larger of the two items (9) and (16); if one tape channel is used, multiply item (9) by 2 and enter their product or item (16), whichever is larger. _____
18. Divide (1e) by (2c) to obtain the number of Phase I output sequences. _____
19. Using (18) extract from Table IV the number of Phase II passes. _____
20. Multiply (17) by (19) to obtain the total Phase II time. _____
21. Extract from Table III:
 - a. Phase III basic process time. _____
 - b. Phase III time for each additional comparison. _____
22. Multiply (1f) by (21b). _____
23. Enter the time for additional Phase III programming. _____
24. If Form 3 records are being sorted, enter (14). _____
25. If hash total option 2, 3, 4, 5, 6, or 7 is specified, enter 0.21. _____
26. Enter the Phase III process time: for a 7070 the sum of items (21a), (22), (23), (24), and (25); for a 7074, the sum of these items divided by 5. _____


27. Using (1b) and (1d), extract the Phase III output tape time from Table V, VI, VII, or VIII. _____
28. Enter the total Phase III time: if two tape channels are used, enter the largest of the three items (9), (26), and (27); if one tape channel is used, add (9) and (27) and enter their sum or (26), whichever is greater. _____
29. Add (10), (20), and (28). _____
30. Multiply (1e) by (29) and divide the result by 1000. _____
31. Divide (30) by 60 to obtain the Sort 90 running time in minutes, exclusive of rewind time. _____
32. Add 2 to item (19) to obtain the required number of rewinds. _____
33. Enter the total rewind time in minutes: if rewinds of full reels are expected, multiply (32) by 0.9 (729 IV) or by 1.2 (729 II); if rewinds of partial reels are expected, multiply (31) by 0.1. _____
34. Add (31) and (33) to obtain the total Sort 90 running time in minutes. _____

APPENDIX III

NON-ENDING SORT CONDITION

The non-ending sort described on page 13, which is possible when an attempt is made to sort more than the allowed maximum number of records, arises because of the procedures by which writing and subsequent merging of sequences of records are accomplished in Phase II of Sort 90. When in a Phase II pass a new output sequence is detected, the tape used for writing the sequence is a tape other than the one used for writing the just-previous output sequence; the available output tapes are actually used in rotating fashion. When the end of a tape is detected, that tape of course becomes unavailable for further writing in the current Phase II pass.

It is very possible that there could be built up after several Phase II passes a small number of very long sequences and a larger number of relatively small sequences. With a three-way merge, a late Phase II pass could produce three sequences as follows:



Sequence 1 occupies all of the first and second reels and the beginning of the third. When the next Phase II pass begins and these tape reels are used for input, the first new output sequence will be simply Sequence 1 again, and the arrangement of sequences on the three new output tapes will be identical to the above picture. No merging has been performed and passes could be executed indefinitely without improving the situation.

If, however, all the records being sorted could be contained on two tape reels, any one sequence regardless of length could not occupy more than two reels, and the third reel would always begin with another sequence, so that merging can continue.

If any Phase II pass (M-1) output reels are filled, where M is the order of merge, a message to this effect will be typed. At the end of each Phase II pass, the program tests to see if (1) any reduction has been made in the number of Phase II passes remaining to be executed and if (2) the number of sequences written by this pass is less than the number written by the previous pass. If the answers are negative for the first condition and positive for the second, at least one output sequence occupies more than one reel of tape, and the program can continue operation with a lower effective order of merge. If, however, the answers are negative for both conditions, the non-ending situation has developed, and the sorting attempt should be terminated.

INDEX

When more than one page reference for a particular subject is listed, the page number underscored indicates the major reference.

- A - units 8
- Additional Files
 - Modification I-9 61
 - Modification III-5 70
- Area Allocation
 - Modification I-4 57
 - Phase I 19
- Arrangement, Optimal 16
- Assignment Routine 4, 51
 - Phase I Exits 56
 - Phase II Exits 63
 - Phase III Exits 66
- B - units 8
- Block
 - Modification I-5 58
 - Size 12
- Calculations, Timing 76
- Channels, Tape 10
- Check, Sequence 5
- Checkpoint 20
 - Modification II-6 65
 - Modification III-7 72
 - Modifications 52
- Collating Sequence 17
- Control Cards 27
 - Checking 27
 - Input File Label 32
 - Input Unit 27
 - Listing 27
 - Modification I-1 55
 - Output File Label 33
 - Reading 27
- Control Data, Sort 15
- Count
 - Record 18
 - Reel 14
- Creation Date 25
- Data Sort Control 15
- Deletions
 - Modification I-6 58
 - Modification III-2 66
- DTF Modifications
 - Phase I 59
 - Phase III 69
- Dump Tape 20
 - Tape Unit 11
- Editing Modifications 52
 - Phase I 58
 - Phase III 66
- EOF Trailer Label 14
- Features of Sort 90 1
- Field, Sort Control Data 15
- Files
 - EOF Trailer Label 14
 - Input 12
- Glossary 74
- Halts see: Programmed Halts
- Hash Totals 18, 50
- Header Label
 - Dump Tape 20
 - Output Tape 25
- I - units 8
- Input File 12
 - Block Size 12
 - EOF Trailer Label 14
 - Label Control Card 32
 - Label Modification I-6 60
 - Number of Records 13
 - Number of Reels 13
 - Record Size 12
 - Record Type 12
 - Tape Units 8
- Internal Latches 52
- Internal Sort 4
- Interruption 24
- Label
 - Dump Tape 20
 - EOF Trailer, Input 14
 - Hash Totals 18
 - Input Control Card 32
 - Input Modification 60
 - Output 25
 - Output Control Card 33
 - Work Tapes 22
- Listing, Control Cards 27
- Main Program 4, 51
- Messages 34
 - Phase I 34
 - Phase II 35
 - Phase III 36
- Modifications 50
 - Checkpoint 52
 - Editing 52
 - Internal Latches 52
 - Patching 50
 - Phase I 55
 - Phase II 62
 - Phase III 66
 - Placement of 51
 - Restart 52
 - Types of 50
- Natural Sequencing 4

- Non-ending Sort 13, 85
- Number of
 - Records 13
 - Reels 13
 - Tape Units 5
- Operation 23
 - Actions during 23
 - Console switches and Dials 25
 - Initiation 23
 - Interruption 24
 - Restart 24
- Optimal Arrangement 16
- Order of Merge 4
- Output File
 - Header Label 25
 - Label Control Card 33
 - Tape Units 10
- Patching, Modifications 50
- Phase I
 - Area Allocation 19
 - Messages 34
 - Modifications 55
 - Program Description 4
 - Programmed Halts 37
 - Record Count 18
- Phase II
 - Messages 35
 - Modifications 62
 - Order of Merge 4
 - Program Description 5
 - Programmed Halts 45
 - Record Count 18
 - Summarization 64
- Phase III
 - Messages 36
 - Modifications 66
 - Program Description 5
 - Programmed Halts 47
 - Record Count 18
 - Summarization 66
- Placement of Modifications 51
- Program Tape Unit 11
- Programmed Halts 37
 - Phase I 37
 - Phase II 45
 - Phase III 47
- Reading Control Cards 27
- Record
 - Count 18
 - Number of 13
 - Size 12
 - Type 12
 - Unreadable 20
- Reel Count 14
- Reels, Input 13
- Restart 20
 - Modifications 52
 - Operation 24
- Retention Cycle 25
- Running Program 4
- Segment, Control Data 15
- Segment Marks 20
- Sequence
 - Check 5
 - Collating 17
 - Natural 4
- Size
 - Block 12
 - Record 12
- Sort Control Data 15
 - Collating Sequence 17
 - Data Fields 15
 - Optimal Arrangement 16
 - Segments 15
- Sorting Times 76
 - Step-by-Step Procedure 82
 - Timing Calculations 76
- Summarization 53
 - Modification II-4 64
 - Modification III-2 66
- Tape
 - Block, Modification I-5 58
 - Channels 10
 - Dump 20
 - Work 22
- Tape Labels
 - Input, Modification I-6 60
 - Modification III-4 70
 - Work Tapes 22
- Tape Units
 - Dump 11
 - Input 8
 - Output 10
 - Program 11
 - Total Number 5
- Timing 76
- Total Number of Tape Units 5
- Totals, Hash 18
- Trailer Label
 - EOF, Input 14
 - Hash Totals 18
- Types of
 - Modifications 50
 - Records 12
- Units, Tape see: Tape Units
- Unreadable Records 20
- Work Tapes 8
 - Labels 22

IBM 7070 PUBLICATIONS

The following IBM 7070 Systems literature has been published as of the date of this manual:

GENERAL INFORMATION MANUALS

<u>Form Number</u>	<u>Title</u>
F22-6517	Introduction to IBM Data Processing Systems
F28-8043	IBM Commercial Translator
F28-8053	The COBOL Translator

REFERENCE MANUALS

C28-6078-1	7070 Basic Autocoder
C28-6090	Simulation of the IBM 650 on the IBM 7070
C28-6099	7070 Basic FORTRAN
C28-6102	7070/7074 Four-Tape Autocoder
C28-6110	7070/7074 Utility Programs
C28-6111	7070/7074 Sort 90

BULLETINS

J20-6067	Uses of Index Words in the IBM 7070
J28-6032-1	IBM 7070 Autocoder
J28-6033-1	IBM 7070 Input/Output Control System
J28-6040	IBM 7070 Sort 90 and Merge 91 Specifications
J28-6041-1	Assembly and Condensing of 7070 Basic Autocoder Programs on the IBM 650
J28-6042-1	Simulation of the IBM 7070 on the IBM 704 and the IBM 7090
J28-6045	IBM 7070 FORTRAN
J28-6047-1	IBM 7070 SPOOL System
J28-6049	IBM 7070 Report Program Generator
J28-6053	Additions to the IBM 7070 Autocoder: Writing Macro Generators for the IBM 7070 Autocoder
J28-6069	IBM 7070 Sort 90 and Merge 91: Timing Estimates; Modifications
J28-6070	Effects of Increased 7070 Tape Capacity and 7070 Tape-Oriented System on IBM Programs
J28-6083	Machine Requirements for IBM Programs and Programming Systems
J28-6095	IBM 7070 Utility Control Panels
J28-6105	IBM 7070/7074 Compiler Systems: Operating Procedure
J28-8023	Machine Optimal Approximations
J28-8026	Specifications for Submittal and Processing of Social Security Tape Reports

The following IBM 7070 Machine literature has been published as of the date of this manual:

GENERAL INFORMATION MANUAL

D22-7004-2	7070 Data Processing System
------------	-----------------------------

REFERENCE MANUALS

A22-6502-1	IBM 700-7000 Series Auxiliary Operations
A22-7003-2	7070 Data Processing System

BULLETIN

G22-6545	General Information (7074 Data Processing System)
----------	---

IBM

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