

IBM System/3 Disk Sort Reference Manual

- Disk Sort Program Numbers:**
- 5702-SM1 (Models 8 and 10)
 - 5703-SM1 (Models 4 and 6)
 - 5704-SM1 (Model 15)
 - 5704-SM9 (Model 15)
 - 5705-SM1 (Model 12)

- CCP/Disk Sort Program Numbers:**
- 5703-SM2 (Model 4)
 - 5704-SM7 (Model 15)
 - 5799-ATF (Models 8 and 10)
 - 5799-ATG (Model 12)
 - 5799-ATH (Model 15)

Program Product

SC21-7522-8
File No. S3-33

Ninth Edition (April 1978)

This is a major revision of, and obsoletes, SC21-7522-7 and technical newsletters SN21-5405, SN21-5511, and SN21-5495. Because the changes and additions are extensive, this manual should be reviewed in its entirety.

This edition applies to the following IBM System/3 programs and to all subsequent versions and modifications until otherwise indicated in new editions or technical newsletters.

Version	Modification	Program Number	Program Name	Model
16	0	5703-SM2	CCP/Disk Sort Program	4
16	0	5799-ATF	CCP/Disk Sort Program	8, 10
5	0	5799-ATG	CCP/Disk Sort Program	12
7	0	5799-ATH	CCP/Disk Sort Program	15
4	0	5704-SM7	CCP/Disk Sort Program	15D
16	0	5703-SM1	Disk Sort Program	6
16	0	5702-SM1	Disk Sort Program	8, 10
5	0	5705-SM1	Disk Sort Program	12
7	0	5704-SM1	Disk Sort Program	15
4	0	5704-SM9	Disk Sort Program	15D

Changes are periodically made to the information herein; these changes will be reported in technical newsletters or in new editions of this publication.

Use this publication only for the purposes stated in the *Preface*.

Publications are not stocked at the address below. Requests for copies of IBM publications and for technical information about the system should be made to your IBM representative or to the branch office serving your locality.

This publication could contain technical inaccuracies or typographical errors. Address comments concerning the contents of this publication to IBM Corporation, Publications, Department 245, Rochester, Minnesota 55901. IBM may use and distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.



Technical Newsletter

This Newsletter No.	SN21-5651
Date	29 September 1978
Base Publication No.	SC21-7522
File No.	S3-33
Previous Newsletters	None

IBM System/3 Disk Sort Reference Manual

© IBM Corp. 1970, 1971, 1972, 1973, 1975, 1976, 1978

This technical newsletter applies to the current versions and modifications of the applicable IBM System/3 programs listed in the edition notice, and provides replacement pages for the subject publication. These replacement pages remain in effect for subsequent versions and modifications unless specifically altered. Pages to be inserted and/or removed are:

Cover, Edition Notice
13, 14
21, 22
25, 26
91 through 110
141, 142
149 through 152

Changes to text and illustrations are indicated by a vertical line at the left of the change.

Summary of Amendments

- Addition of Release/Task Chaining capabilities.
- Miscellaneous technical and editorial changes.

Note: Please file this cover letter at the back of the manual to provide a record of changes.

IBM Corporation, Publications, Department 245, Rochester, Minnesota 55901

© IBM Corp. 1978

Printed in U.S.A.

This Newsletter No. SN21-5704

Date 28 March 1980

Base Publication No. SC21-7522-8

File No. S3-33

Previous Newsletters SN21-5651

**IBM System/3
Disk Sort
Reference Manual**

© IBM Corp. 1970, 1971, 1972, 1973, 1975, 1976, 1978

This technical newsletter applies to the current versions and modifications of the applicable IBM System/3 programs listed in the edition notice, and provides replacement pages for the subject publication. These replacement pages remain in effect for subsequent versions and modifications unless specifically altered. Pages to be inserted and/or removed are:

cover, edition notice
7 through 10
13, 14
14.1, 14.2 (added to accommodate
new and moved text)
23, 24
24.1, 24.2 (added to accommodate
new and moved text)
45, 46
69, 70
107, 108
123, 124
147, 148
157, 158
161 through 164

Changes to text and illustrations are indicated by a vertical line at the left of the change.

Summary of Amendments

Miscellaneous technical changes

Note: Please file this cover letter at the back of the manual to provide a record of changes.

This manual helps the System/3 programmer fill out the sequence specifications needed in order to run:

- The System/3 Disk Sort program on System/3 Models 4, 6, 8, 10, 12, and 15
- The System/3 CCP/Disk Sort program on System/3 Models 4, 8, 10, 12, and 15

Two coding sheets are discussed in this manual:

- *Sequence Specifications*, GX21-9089
- *Translation Table and Alternate Collating Sequence Coding Sheet*, GX21-9096

The *How to Use This Manual* section in Chapter 1 explains the organization of this book.

For ease of illustration, many of the examples in this book use card-like figures to represent records. This does not imply that a card device must be used for input or output in these situations. Any of several input/output devices might be used, depending on which System/3 model and configuration you are using.

IBM System/3 5448 Disk Storage Drive

The IBM System/3 5448 Disk Storage Drive on System/3 Models 8 and 10 uses the same program product support as the IBM 5445 Disk Storage. However, a separate system control program feature is required for the 5448. In general, references to 5445 in this manual also apply to 5448. For specific information about 5448 operating characteristics and programming support, see *IBM System/3 5448 Disk Storage Drive Program Reference Manual*, GC21-5168.

System/3 Model 4 users should note that:

- System/3 Model 4 is supported by System/3 Model 6 System Control Programming and Program Products. The facilities described in this publication for the Model 6 are also applicable to the Model 4, although the Model 4 is not referenced.
- Not all devices and features available on the Model 6 are available on the Model 4. Therefore, Model 4 users should be familiar with the contents of the *IBM System/3 Model 4 Introduction*, GC21-5146.
- All references to the 5444 disks on systems without 3340 Direct Access Storage Facility also apply to the 5447 disks unless otherwise noted.

System/3 Model 8 users should note that:

- The System/3 Model 8 is supported by System/3 Model 10 Disk System System Control Programming and Program Products. The facilities described in this publication for the Model 10 are also applicable to the Model 8, although the Model 8 is not referenced.
- Not all devices and features available on the Model 10 are available on the Model 8. Therefore, Model 8 users should be familiar with the contents of the *IBM System/3 Model 8 Introduction*, GC21-5114.

System/3 Model 10 Disk System users should note that:

- In this publication the System/3 Model 10 Disk System is termed the Model 10. Without exception, all references to Model 10 apply to the Model 10 Disk System.

System/3 CCP/Disk Sort Program for Models 8, 10, 12, and 15

Appendix I points out the differences between the CCP/Disk Sort Programs for Models 8, 10, 12, and 15 and the corresponding Disk Sort Program for these systems.

System/3 Models 12 and 15 users should note that:

- The IBM 3340 Direct Access Storage Facility attaches to the System/3 Model 12 and to the System/3 Models 15B, 15C, and 15D.
- The IBM 3344 Direct Access Storage attaches to the System/3 Model 15D.
- Certain areas on the 3340 and 3344 disks are treated as 5444 disks. These areas are known as *5444 simulation areas* and are used for the program libraries and can be used for data files. The remainder of the disk space is known as the *main data area* and can only be used for data files.
- References in this manual to the 5444, 5445, and 3340 are to be interpreted according to which disk storage device(s) is(are) attached to your system. The following table is used to determine the meaning of the reference:

	Models 8 and 10	Model 15A	Models 12, 15B, 15C	Model 15D
Reference	Meaning	Meaning	Meaning	Meaning
5444	5444 Disk Storage Drive	5444 Disk Storage Drive	5444 simulation area on 3340	5444 simulation area on 3340 or 3344
5445	5448 Disk Storage Drive or (Model 10) 5445 Disk Storage	5445 Disk Storage	Main data area on 3340	Main data area on 3340 or 3344
3340	Not applicable	Not applicable	Main data area on 3340	Main data area on 3340 or 3344

- All references to the System/3 Model 15 in this manual apply to all models of Model 15 as specified below unless otherwise indicated:
 - For Disk Sort, program number 5704-SM1 for Models 15A, 15B, and 15C; and program number 5704-SM9 for Model 15D.
 - For CCP/Disk Sort, program number 5704-SM7 for Model 15D.

For further information, see the appropriate publication listed under *Related Publications*; for more information about the manuals or the coding sheets, contact your nearest IBM branch office.

RELATED PUBLICATIONS

Type of IBM System/3 Publication	Order Number					
	Model 4	Model 6	Model 8	Model 10	Model 12	Model 15
Introduction	GC21-5146	GA21-9122	GC21-5114	GC21-7510	GC21-5116	GC21-5094
System Control Program Reference			GC21-7512		GC21-5130	GC21-5077 (5704-SC1) GC21-5162 (5704-SC2)
OCL and Disk Utilities	GC21-7516					
Components Reference	GA34-0001		GA21-9236			
Disk Concepts and Planning Guide	GC21-7571					
Operator's Guide	GC21-5149	GC21-7501	GC21-7508		GC21-5144	GC21-5075
Messages/Halt Guide	GC21-7541		GC21-7540		GC21-5145	GC21-5076
Communications Control Program Terminal Operator's Guide	GC21-7580		GC21-7580			
CCP Programmer's Reference Manual	GC21-5150		GC21-7579			
<i>Program Logic Manuals:</i>						
Disk Sort and CCP/Disk Sort PLM	LY21-0517					
System Control Program PLM	SY21-0502				SY21-0046	SY21-0033
Scheduler PLM						SY21-0035
Data Management PLM	SY21-0512					SY21-0034

Contents

CHAPTER 1. INTRODUCTION	1	CHAPTER 4. RECORD TYPE SPECIFICATIONS	26
How to Use This Manual	1	Column Summary	26
Important Information for Users of System/3		Column Descriptions	27
CCP/Disk Sort	1	Columns 1-2 (Page Number) and 3-5 (Line Number)	27
System Configuration	2	Column 6 (Line Type)	27
CHAPTER 2. OVERVIEW OF DISK SORT	3	Include Sets	28
The Three Types of Sort: Addrout, Tag Along, and		Omit Sets	29
Summary Tag Along	5	Column 7 (Continuation or Comments)	30
Running the Disk Sort or CCP/Disk Sort Program	6	Column 8 (C/Z/D/P/U)	30
Sequence Specifications Sheet	7	Columns 9-16 (Factor 1 Location)	34
Order of Sequence Specifications	8	Columns 17-18 (Relation)	34
How to Supply Sequence Specifications	8	Column 19 (Field or Constant)	35
OCL Statements	9	Columns 20-27 (Factor 2 Field)	35
File Considerations	10	Columns 20-39 (Factor 2 Constant)	35
7-Track Tape Considerations	15	Columns 40-72 (Comments)	36
9-Track Tape Considerations	16	CHAPTER 5. FIELD DESCRIPTION SPECIFICATIONS	37
Restrictions on the Use of Multifile Tapes (Model 15)	16	Column Summary	37
How to Supply OCL Statements for the Model 6	16	Column Description	38
How to Supply OCL Statements for the Model 10	16	Columns 1-2 (Page Number) and 3-5 (Line Number)	38
How to Supply OCL Statements for the Model 12	17	Column 6 (Line Type)	38
How to Supply OCL Statements for the Model 15	17	Column 7 (Field Type or Comments)	38
Timing Considerations	17	Column 8 (C/Z/D/P/U/V)	40
Programming Considerations	18	Columns 7-8	43
Specifications for Tag Along Sort (SORTR)	19	Columns 9-16 (Field Location)	44
Specifications for Addrout Sort (SORTA)	20	Column 17 (Conditionally Forced Character)	45
Specifications for Summary Sort (SORTRS)	21	Column 18 (Forced Character)	45
CHAPTER 3. HEADER SPECIFICATIONS	22	Column 19 (Specifying Replacements for More Than One	
Column Summary	22	Character)	46
Column Descriptions	23	Examples of Using Forced Control Fields	47
Columns 1-2 (page Number) and 3-5 (Line Number)	23	Columns 20-22 (Overflow Field Length Entry)	52
Column 6 (Line Type)	23	Columns 23-39	52
Columns 7-12 (Job)	23	Columns 40-72 (Comments)	52
Columns 13-17 [Largest Sum (in Bytes) of Control Field		CHAPTER 6. SAMPLE DISK SORT JOBS	53
Lengths for Any Record Type]	23	Record Formats for the Inventory File	54
Column 18 (Ascending or Descending Sequence)	23	Record Formats for the Transaction File	54
Column 26 (Collating Sequence)	23	Sample Job 1: Include on Field Relationship	55
Column 27 (Print Option)	24	Header Specifications	56
Column 28 (Output Option for Tag Along Sorts)	24	Record Type Specifications	56
Columns 29-32 (Output Record Length for Tag Along Sorts)	24	Field Specifications	57
Column 34 (Nonverify Option)	24	Sample Job 2: Addrout Sort	58
Columns 37-38 (Storage Size) — CCP/Disk Sort Model 15		Header Specifications	58
5704-SM7 and 5799-ATH	25	Record Type Specifications	59
Column 39 (Record Length Indicator) — CCP/Disk Sort Only	25	Field Specifications	59
Columns 40-43 (Record Length) — CCP/Disk Sort Only	25	Sample Job 3: And Logic—Opposite Control Fields	60
Columns 40-72 (Job Description) or Columns 44-72 (Job		Header Specifications	61
Description — CCP/Disk Sort Only)	25	Record Type Specifications	61
Columns 75-80 (Program Identification/Program Name)	25	Field Specifications	62
		Sample Job 4: Multiple Include Sets with Force	63
		Header Specifications	65
		Record Type Specifications	65
		Field Specifications	65
		Sample Job 5: Summary Sort—Eliminating Duplicates	67
		Header Specifications	68
		Record Type Specifications	68
		Field Specifications	68

Sample Job 6: Summary Sort—Insertion of Character		APPENDIX D. TIMING CONSIDERATIONS	110
Data—FDV	69	Main Storage Size	110
Header Specifications	70	Number of Records to be Sorted	110
Record Type Specifications	70	Record Size	110
Field Specifications	70	Number of Sequence Specifications	111
Sample Job 7: Negative Values Possible	72	Alternate Collating Sequence	111
Header Specifications	72	Disk Location of the Files	111
Record Type Specifications	72	Order of the Control Fields in the Input File	113
Field Specifications	73	Use of the Automatic Work File Allocation Function	113
Sample Job 8: Multiple Input Files	74	Nonverify Option	113
Header Specifications	76	Sample Disk Sort Timings	113
Record Type Specifications	76	Model 6 and Model 10 Timings Only	113
Field Specifications	77	Model 15 Timings	113
Sample Job 9: Record Selection on Binary Values	78	Sample Model 6 and Model 10 Timings	114
Type 1 Include Sequence	78	Sample Model 15 Timings	115
Type 2 Include Sequence	78	Sample Model 12 Timings	115
Sample Sort Job 9 Description	79		
Header Specifications	82	APPENDIX E. CALCULATING SIZE OF FILES	116
Record Type Specifications	82	Input File	116
Field Specifications	82	Output File	116
		Work File	116
CHAPTER 7. INTEGER (BINARY) AND REAL NUMBER		Work File Size Factor	118
SORTING	83	Example	119
Sample Job 1: Sorting on a Control Field that Contains Either		5444 Disk Files	119
a Positive Integer or a Positive Real Number	83	5445 Disk Files	120
Header Specification	84	3340 Disk Files (Model 12 and Model 15)	120
Record Type Specifications	84		
Field Specifications	84	APPENDIX F. MESSAGES	121
Sample Job 2: Sorting on a Control Field that Contains Only			
a Positive or Negative Integer	85	APPENDIX G. INSTRUCTIONS FOR EXECUTING THE DISK	
Header Specification	86	SORT SAMPLE PROGRAM	146
Record Type Specifications	87	Sample Program for Model 6	146
Field Specifications	87	Sample Program for Model 8, Model 10, Model 12, and	
Sample Job 3: Sorting on a Control Field that Contains Only		Model 15	147
a Positive or Negative Real Number	88		
Header Specification	89	APPENDIX H. CCP/DISK SORT SPECIFICATIONS	
Record Type Specifications	90	DIFFERENCES	148
Field Specifications	90	CCP/Disk Sort Overview	148
		File Statement Considerations	148
APPENDIX A. CONVERSATIONAL OCL FOR MODEL 6	91	Multiple Input Files (Model 15 only)	149
Coding OCL Cycles	91	Output File Over Input File Considerations for Model 4	149
The Modify Statement	91	Work Record Length	149
The Load Cycle	92	CCP/Disk Sort Object Program Generation Job Stream	
The Build Cycle	95	Example	150
The Call Cycle	99	CCP Assignment Set Statements Example	152
		CCP/Disk Sort Execution (Batch Mode) Job Stream	
APPENDIX B. SUPPLYING SEQUENCE SPECIFICATIONS	100	Example	153
Interrelationship of Conversational OCL and Sequence			
Specifications (Model 6)	101	APPENDIX I. CCP/DISK SORT PROGRAM FOR SYSTEM/3	
Supplying OCL Statements and Sequence Specifications for the		MODELS 8, 10, 12, 15A, 15B, AND 15C	154
Model 10, the Model 12, and the Model 15	102	CCP/Disk Sort Program (Program Number 5799-ATF)	154
Which OCL and Sequence Specification Method to Use for the		CCP/Disk Sort Program (Program Number 5799-ATG)	154
Model 10, the Model 12, and the Model 15	105	CCP/Disk Sort Program (Program Number 5799-ATH)	154
APPENDIX C. COLLATING SEQUENCE	106	APPENDIX J. GLOSSARY	155
Standard Collating Sequence	107		
Alternate Collating Sequence	107	INDEX	157
Coding Rules	107		
Input Order when ALTSEQ Statements are Used	108		
Programming Considerations	108		
Effect of ALTSEQ Statements on Other Coding	108		
Sample ALTSEQ Statements	108		
Programming Aids	108		

This publication is designed to help you, the programmer, use System/3 Disk Sort and System/3 CCP/Disk Sort. The CCP/Disk Sort Program is used with the Communications Control Program on the System/3 Model 4 with Program Number 5703-SM2 and Model 15 with Program Number 5704-SM7.

HOW TO USE THIS MANUAL

Your need for reference material varies with programming experience and familiarity with a particular system. This manual contains four different levels of information:

1. *Overview of Sequence Specifications* shows which columns on the sequence specifications sheet you must consider when you want to sort a disk or tape file.
2. *Column Summaries* list the possible entries for each column on the sequence specification sheet.
3. *Column Descriptions* explain the possible entries for each column in greater detail.
4. *Sample Jobs* show how to fill out the sequence specification sheets for typical jobs.¹

Here is how to use each level:

- Use the *Overview* to see which columns on the sequence specifications sheet you must consider when you want to sort a disk or tape file.
- Use the *Column Summaries* for a quick recall of the possible entries for each column on the sequence specifications sheet.
- Use the *Column Descriptions* when you need a detailed explanation of the entries for a particular column.
- Use the *Sample Jobs* to review coding techniques for actual jobs.¹

¹See Appendix H for CCP/Disk Sort differences.

A column summary and column descriptions are provided for each type of specification:

- Header
- Record type
- Field description

Important Information for Users of System/3 CCP/Disk Sort

System/3 CCP/Disk Sort users should review Appendix H prior to coding your program for the following reasons:

- If you are knowledgeable with System/3 Disk Sort programming and System/3 CCP, Appendix H illustrates the CCP/Disk Sort differences your coding must reflect.
- Appendix H provides a CCP/Disk Sort overview. For more details, refer to the chapter discussions.
- If you are a first time user of System/3, System/3 Disk Sort, or System/3 CCP/Disk Sort, Appendix H provides an orientation to System/3 CCP/Disk Sort for your System/3 Model 4 (5703-SM2) or Model 15 with Program Number 5704-SM7.

Additional information on System/3 CCP/Disk Sort Programs for System/3 Models 8 and 10 (5799-ATF), System/3 Model 12 (5799-ATG), and System/3 Models 15A, 15B, and 15C (5799-ATH) is provided in Appendix I.


SYSTEM CONFIGURATION

For information concerning the minimum system configuration for disk sort and additional devices supported, see one of the following publications, as appropriate for your System/3 model:


- *IBM System/3 Models 4, 6, 8, 10, and 12 System Generation Reference Manual*, GC21-5126
- *IBM System/3 Model 15 System Generation Reference Manual*, GC21-7616
- *IBM System/3 Model 6 Introduction*, GA21-9122
- *IBM System/3 Model 8 Introduction*, GC21-5114
- *IBM System/3 Model 12 Introduction*, GC21-5116
- *IBM System/3 Model 15 Introduction*, GC21-5094

With the IBM Disk Sort or CCP/Disk Sort (also see Appendix H) Program you can:


- Rearrange the records in a disk or tape file:



- Drop records from a disk or tape file:




- Reformat the records in a disk or tape file:



Here is how the Disk Sort or CCP/Disk Sort Program works with the records in your file:

1. Disk sort reads a record from the input file:




Note: Control fields are the fields you want disk sort to use to sort the records. See *Column 7* under *Field Description Specifications* for a detailed discussion of control fields.

2. Disk sort checks your sequence specifications to see whether the record is one you want to sort (often you may not want to sort all the records in the file).
3. If the record *is* one you want to sort, disk sort builds a work record, formatting it according to your sequence specifications (the format of the work record is important because it controls the format of the output record).

Assume that your sequence specifications say:

- a. Put the contents of positions 27-30 in the input record into positions 1-4 of the work record.
- b. Put the contents of positions 1-5 in the input record into positions 5-9 of the work record.
- c. Put the contents of positions 6-26 in the input record into positions 10-30 of the work record.

The work record disk sort builds would look like this:



See *Column 7, Field Description Specifications* for a more detailed discussion of the work record.

If you are not dropping control fields, they will always precede data fields in the work and output records (see *Column 28-Header Specifications* for a discussion of dropping control fields).

4. Disk sort writes all the records you want to sort into the work file.
5. Disk sort checks your sequence specifications to see how you want to arrange the records in the output file (ascending or descending order by control fields).
6. Disk sort writes the records in the output file in the order you have specified. The only exception is when you want to drop control fields from the output record.

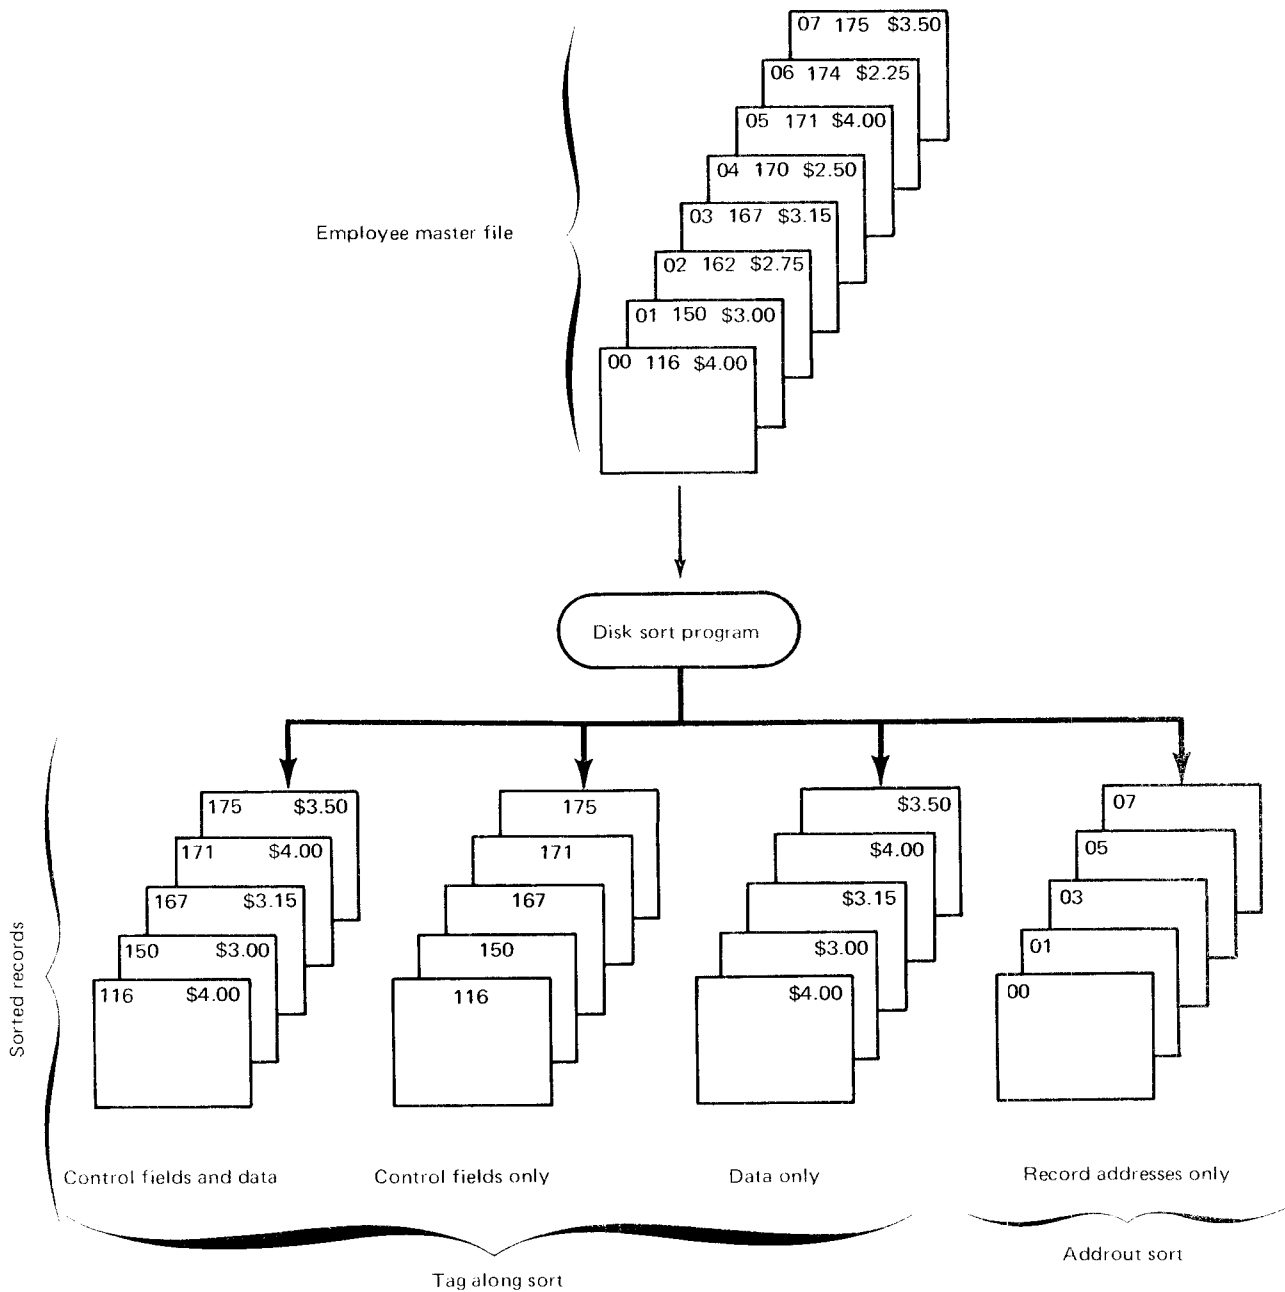
THE THREE TYPES OF SORT: ADDRROUT, TAG ALONG, AND SUMMARY TAG ALONG

There are three types of sort jobs: addrout (address out), tag along (data fields can *tag along* with control fields in the sorted records), and summary tag along (data is summarized in the sorted records).

The output from an addrout sort job consists of 3-byte binary relative record numbers of the records in the input file.

The output for a tag along sort is a file of sorted records. The sorted records can contain:

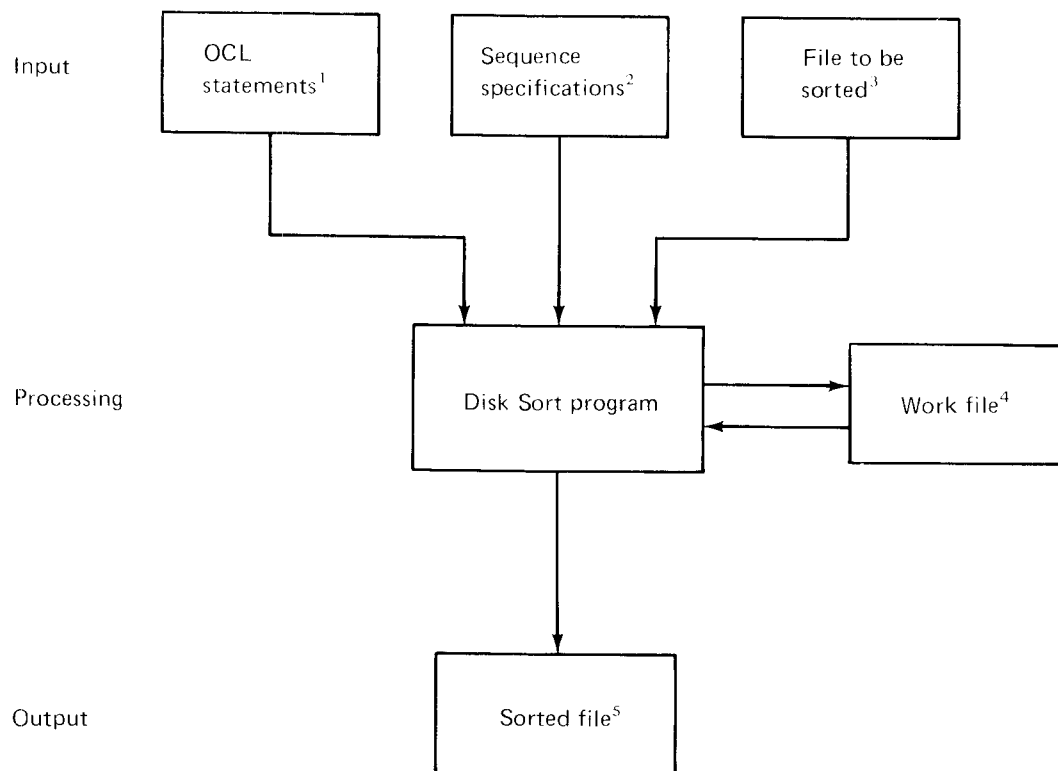
- Control fields and data
- Control fields only
- Data only



The output for a summary tag along sort is a file of sorted records. The sorted records can contain:

- Control fields, data fields, and summary data
- Control fields only
- Data fields and summary data
- Summary data fields only

RUNNING THE DISK SORT OR CCP/DISK SORT PROGRAM



¹ *OCL statements* are your instructions to the computer.

² *Sequence specifications* are your instructions to the disk sort program.

³ *File to be sorted* can be an indexed, sequential, or direct disk file, or a sequential tape file (or card or 3741 diskette file for Mode 15).

⁴ *Work file* is an area on disk where disk sort writes all the records you want to sort. The program can then work with the records there without disturbing the input file.

⁵ *Sorted file* can contain:

- a. The relative record numbers of the records in the file
- b. Part or all of the records themselves (forming a sequential file)

SEQUENCE SPECIFICATIONS SHEET

The sequence specifications sheet contains three different kinds of specifications: header, record type, and field description. The sample specifications sheet below shows the location and purpose of all three.

IBM		SEQUENCE SPECIFICATIONS	
		Header	
Here you specify the type of sort job you are going to run.		Header specifications	
		Record Type	
Here you specify which of the file's records you want to sort.		Record type specifications	
		Field	
Here you specify how you want the records sorted.		Field description specifications	

Order of Sequence Specifications

The normal order of sequence specifications is:

1. Header line
2. Record type lines
3. Field description lines

This order can vary, however, depending on how many records you want to sort and the format of those records (format refers to the locations, lengths, and types of fields in a record). If, for example, you want to sort all the records in a file — and they all have the same format — you do not have to fill out record type specifications (this is often referred to as implied include-all).

Number of Records to Be Sorted	Format of Records to Be Sorted	Order of Sequence Specifications
All the records in the file (implied include-all)	All the same	<ol style="list-style-type: none"> 1. Header line 2. Field description lines
Some of the records in the file	All the same	<ol style="list-style-type: none"> 1. Header line 2. Record type lines 3. Field description lines
All or some of the records in the file	Several different formats	<ol style="list-style-type: none"> 1. Header line 2. Record type lines 3. Field description lines 4. Record type lines 5. Field description lines 6. Record type lines 7. Field description lines <p style="margin-left: 20px;"> { For first type of record format { For second type of record format } One set for each additional type of record format¹ </p>
¹ This does not mean that the records in the file must be grouped by format type. The sequence specifications must be grouped (one set of record type and field description lines for each format type). The records themselves can be in any order.		

For detailed information on how to include or omit certain records when you are sorting a file, see *Column 6* under *Record Type Specifications*.

How to Supply Sequence Specifications

You can supply the sequence specifications from:

- The device you are using to supply the OCL statements
- An OCL procedure containing the LOAD, FILE, and RUN statements for the job (the OCL procedure may or may not contain other jobs)
- A source library module

Appendix B summarizes:

- The advantages and disadvantages of each of these methods
- The interrelationship of the three methods of supplying sequence specifications and three of the OCL cycles (LOAD, BUILD, and CALL)

Note: Regardless of the method used to read the sort specification statements, the system input device is dedicated to the partition in which the sort is running until all such statements are read.

OCL STATEMENTS

Before your system can run a sort job, it needs certain information about the disk sort program, the sequence specifications for the job, and the files involved in the job (every sort job uses a minimum of three files: an input file, a work file, and an output file). You supply this information in OCL statements. Here is the information you must supply:

About the Sequence Specifications

1. How you are going to enter the sequence specifications (for example, keyboard for the Model 6, MFCU for the Model 10, 1442 or MFCU for the Model 12, or display station or MFCU for Model 15).
2. Where you want any computer-to-operator messages printed.

About the Program

1. Name of the program you want to run (\$DSORT).
2. Disk drive location of the disk sort program.

About the Input File

1. Name of the file you want to sort. The FILE statement name must be INPUT (or INPUT1, . . . , INPUT8 for Model 12 or Model 15). However, if the file name is not INPUT (or INPUT1, . . . , INPUT8 for Model 12 or Model 15), use the LABEL keyword to identify the actual file name.
2. Unit assignment of the file you want to sort:
 - a. *All Models Except Model 4:* Removable disk 1 or 2 (R1 or R2) or fixed disk 1 or 2 (F1 or F2) for the 5444
Model 4: Removable disk R1 or fixed disk R2, F1, or F2 for the 5447
 - b. *Model 10:* Drive 1 or 2 (D1 or D2) for the 5445 and the 5448.
 - c. *Model 12:* Drive 1 or 2 (D1 or D2) for the 3340 main data area
 - d. *Model 15:* Drive 1, 2, 3, or 4 (D1, D2, D3, or D4) for the 5445 or 3340
 - e. *Model 15D:* Drive 3 [D31 (D3), D32, D33, and D34] or Drive 4 [D41 (D4), D42, D43, and D44] for the 3344.
 - f. Drive 1, 2, 3, or 4 (T1, T2, T3, or T4) for tape
 - g. Also supported on the Model 15 are MFCU, MFCM, 1442, 2501, and 3741
 - h. If the file is spooled to disk, sort uses the file as if it were coming from the device of original entry.

3. Name of the disk pack or tape that contains the file.
4. If you use tape, specify UNIT, BLKL, and RECL.
5. If you use the 3741, specify RECL (Model 15 only).

About the Work File¹

1. Name of the file you are going to use as a work area. The FILE statement name must be WORK, however, if the name of your file is not WORK, use the LABEL keyword to indicate the actual name of the file.
2. Disk drive location (unit) of the work file:
 - a. *All Models Except Model 4:* Removable disk 1 or 2 (R1 or R2) or fixed disk 1 or 2 (F1 or F2) for the 5444
Model 4: Removable disk R1 or fixed disk R2, F1, or F2 for the 5447
 - b. *Model 10:* Drive 1 or 2 (D1 or D2) for the 5445 and the 5448.
 - c. *Model 12:* Drive 1 or 2 (D1 or D2) for the 3340 main data area
 - d. *Model 15:* Drive 1, 2, 3, or 4 (D1, D2, D3, or D4) for the 5445 and 3340
 - e. *Model 15D:* Drive 3 [D31 (D3), D32, D33, and D34] or Drive 4 [D41 (D4), D42, D43, and D44] for the 3344.
3. Name of the disk pack that contains the file.
4. Number of tracks in the file. (Do not specify number of records.)
5. Status of the file. The work file should be a scratch file since you won't need its information after a sort job.

Note: The work file must be online at all times during the sort run.

¹If the file you want to use is an old (temporary) file, you must also supply the location at which the file begins.

About the Output File¹

1. Name you are going to call the sorted file. The FILE statement name must be OUTPUT; however, if the name of your file is not OUTPUT, use the LABEL keyword to indicate the actual name of the file.
2. Unit assignment of the sorted file:
 - a. *All Models Except Model 4:* Removable disk 1 or 2 (R1 or R2) or fixed disk 1 or 2 (F1 or F2) for the 5444
Model 4: Removable disk R1 or fixed disk R2, F1, or F2 for the 5447
 - b. *Model 10:* Drive 1 or 2 (D1 or D2) for the 5445
 - c. *Model 12:* Drive 1 or 2 (D1 or D2) for the 3340 main data area
 - d. *Model 15:* Drive 1, 2, 3, or 4 (D1, D2, D3, or D4) for the 5445 or 3340
 - e. Drive 1, 2, 3, or 4 (T1, T2, T3, or T4) for tape
 - f. *Model 15D:* Drive 3 [D31 (D3), D32, D33, and D34] or Drive 4 [D41 (D4), D42, D43, and D44] for the 3344.
3. Name of the disk pack or tape on which you are going to put the sorted file.
4. Number of records or tracks in the file.
5. Status of the file. If the file exists before this sort, it cannot have permanent status. If the file is built during the sort, it can have temporary or permanent status.
6. If tape is used, UNIT, BLKL, and RECL must be specified.

Notes:

1. You may not have your output file deferred if it is a split cylinder file.
2. If the END parameter is not used, disk sort defaults to rewinding and unloading tape files.

File Considerations

Input: General

The input file contains records to be sorted. Any file created by System/3 data management can be used as input to disk sort; for example, an RPG II file or a file from a previous disk sort run. The disk file may have been created by any processing method – sequential, indexed, or direct. A tape file must have been created sequentially. Since direct file areas are set to blanks before the file is created, you should test for and omit blank records. You can do this by testing a position that is known never to contain blanks.

¹ If the file you want to use is an old (temporary) file, you must also supply the location at which the file begins.

Input: Model 6 and Model 10

The FILE statement name must be INPUT. However, the actual file need not be called INPUT because the LABEL keyword allows you to identify the file you will be sorting. The input file may be multivolume (online or offline). However, for SORTA jobs, when the input file is processed by an RPG II program with your sorted addrout file, the input file must be online.

The volumes containing the INPUT file must be mounted online *prior* to running disk sort. However, if the input file is offline multivolume, (more than one volume assigned to a removable unit), only the first volume assigned to each unit is required to be mounted. If the input file is not online prior to running, the job will end with message SD 105 (INPUT FILE NOT ONLINE) followed by a CD45 halt (Model 6) or a 25 halt (Model 10).

If your input file is ASCII and ASCII-YES is not specified on the input file card, all data is treated as if it is EBCDIC. EBCDIC and ASCII data cannot be mixed in a sort run. An ASCII input file requires an ASCII output file and an EBCDIC input file requires an EBCDIC output file.

When ASCII is specified for input and output, packed data should not be specified in the sort specifications since ASCII does not support packed data.

Variable-length blocks are not allowed by disk sort tape I/O; the job will end if they are encountered.

Multiple Input: Model 6 and Model 10

Disk sort supports multiple file input to the same extent that System/3 data management supports multiple file input. For more information on this subject, see the multivolume file discussion in the *IBM System/3 Model 10 Disk System Control Programming Reference Manual*, GC21-7512 or the *IBM System/3 Model 6 Operation Control Language and Disk Utility Programs Reference Manual*, GC21-7516.

Restrictions for multiple file input to disk and tape files are:

Disk Multiple File Input: Multiple files can be used as input if they are made to appear as multivolume files with the following restrictions:

1. Labels must all be the same.
2. Record lengths must all be the same.
3. Files reside on different packs.

Tape Multiple File Input (Model 10): Multiple files can be used as input if they are specified as multivolume files with the following restrictions:

1. RECL and BLKL must be the same.
2. Record format must be the same.
3. Unlabeled or nonstandard tape files.
4. Files must have the same attributes (such as, use of a 9-track, 7-track, ASCII, EBCDIC, parity, density, CONVERT, or TRANSLATE).

Note: Standard labeled tape files cannot be used for multiple file input.

Input: Model 12

Input can be from either of the following devices:

- 3340 Direct Access Storage Facility
- 3410/3411 Magnetic Tape Subsystem

Up to eight files can be used as input using any combination of the preceding devices. The file names must be: INPUT or INPUT1 (but not both), INPUT2, INPUT3, INPUT4, INPUT5, INPUT6, INPUT7, and/or INPUT8.

The files will be processed INPUT (or INPUT1) serially to INPUT8 regardless of the order of the OCL FILE statements. The lowest numbered input file will be processed fully first, then the next numbered input file will be processed in its entirety, then the next until all the files have been individually processed in sequence. For a particular disk sort run, the input files do not have to be numbered serially; that is, the file names may be INPUT3, INPUT5, and INPUT8, or a particular run could have just INPUT8 as the input file.

All single-volume disk files must be online prior to running disk sort. However, when using multivolume offline disk files, only the first volume assigned to each unit must be mounted prior to running disk sort. When using 3340 multivolume offline disk files, D1 cannot be used.

Input record length is defined as the record length used for the disk and/or tape files.

Special Considerations: When more than one input file is used, the following must be considered:

- Addrout sort is not allowed.
- Auto allocation of the work file should be used with caution. For best performance, use a work file statement.
- The record lengths of all disk or tape files must agree. This is not true for block lengths. Tape files may have different block lengths.

Data Considerations: When running in a multiple device mode, the data restrictions for any one input device apply to all input devices. For example, if disk input is combined with a 7-track translate input file, packed fields should not be used in the disk sort specifications since a 7-track translate record can contain only the System/3 64-character set.

Input can be mixed EBCDIC and ASCII with either EBCDIC or ASCII output. However, when ASCII is specified, packed data should not be specified in the disk sort specifications since ASCII does not support packed data.

Input: Model 15

Input can be from any of the following devices:

- 5444 Disk Storage Drive
- 5445 Disk Storage
- 3340 Direct Access Storage Facility
- 3344 Direct Access Storage (Model 15D only)
- 3410/3411 Magnetic Tape Subsystem

Note: For the following devices, use a device independent FILE statement.

- 1442 Card Read Punch
- 2501 Card Reader
- 3741 Data Station or 3741 Programmable Work Station directly attached
- 2560 Multi-Function Card Machine (MFCM) or
5424 Multi-Function Card Unit (MFCU)

Up to eight files can be used as input using any combination of the preceding devices. The file names must be: INPUT or INPUT1 (but not both), INPUT2, INPUT3, INPUT4, INPUT5, INPUT6, INPUT7, and/or INPUT8.

The files will be processed INPUT (or INPUT1) serially to INPUT8 regardless of the order of the OCL FILE statements. The lowest numbered input file will be processed fully first, then the next numbered input file will be processed in its entirety, then the next until all the files have been individually processed in sequence. For a particular disk sort run, the input files do not have to be numbered serially; that is, the file names may be INPUT3, INPUT5, and INPUT8, or a particular run could have just INPUT8 as the input file.

Card input files will have a record-length characteristic of the device. If card input files are used with disk or tape files, the card input is truncated or padded on the right with blanks to the record length specified for the disk or tape files.

Single-volume disk files need not be online prior to running disk sort. However, when using multivolume offline disk files, the first volume assigned to each unit must be mounted prior to running disk sort. When using 3340 multivolume offline disk files, D1 cannot be used.

Input record length is defined as the record length used for the disk and/or tape files. If the 3741 and/or cards are the only input, the input record length is the largest specified for the 3741 or 80 or 96 for card input -- whichever is largest.

Special Considerations: When more than one input file is used, the following must be considered:

- Addrout is not allowed.
- Auto allocation of the work file should be used with caution. For best performance, use a work file statement.
- The record lengths of all disk (5444, 5445, or 3340) or tape files must agree. This is not true for block lengths. Tape file may have different block lengths.

Data Considerations: When running in a multiple device mode, the data restrictions for any one input device apply to all input devices. For example, if disk input is combined with 5424 MFCU input, packed fields should not be used in the disk sort specifications since 96-column cards can contain only the System/3 64-character set.

Input can be mixed EBCDIC and ASCII with either EBCDIC or ASCII output. However, when ASCII is specified, packed data should not be specified in the disk sort specifications since ASCII does not support packed data.

Considerations for Writing the Output File Over the Input File: If the output file is to be written over the input file (load to old), the output file will be scratched at the end of job under the following conditions, as though RETAIN-S had been specified for the output FILE statement:

1. A pack containing an input file is not online at the start of the job (deferred mount)
2. RETAIN-S is used on the FILE statement for the *input* file

To avoid this problem, use RETAIN-T for the input FILE statement.

CAUTION

When running disk sort program (5704-SM9) using output over input, take care to ensure that another partition does not share the file to be overlaid during sort execution. If this condition occurs, sort may not be able to reopen the file as an output file at output time. This may result in an unrecoverable halt.

Work File

The work file is generally an area on disk which you have set aside as a work area. The disk sort program uses this area as working storage during the sorting process.

The FILE statement name must be WORK. If the name of the work area is not WORK, use the LABEL keyword to indicate the file you will be using.

A work file should be given scratch status (RETAIN-S), because it is normally created during the sort and has no function afterwards. It is always allocated with a multivolume attribute, even though its FILE statement may indicate single volume.

The work file may be multivolume, but all of the volumes must be online. If either the input or output files are multivolume offline, the removable units containing these multivolume offline files may not be used for the work file.

The work file may be specified on the 5444 or 5445 drives but not on both. An error occurs if work files are allocated on both 5444 and 5445 drives.

When specifying the work file on the 3340 (Model 12 and Model 15), be certain to specify only the 3340 main data area or the simulation area, but not both.

When a new output file is specified and a work FILE statement is given with the same unit as the output file, a halt can occur. The halt is issued when the space allocated for the work file overlaps the area where the new output file is to be located. To avoid the halt, either omit the work FILE statement and let automatic allocate handle the work file or specify locations for both the work file and new output file.

Disk sort work files are opened with multivolume file attributes. If RETAIN-S is used for sort with Model 15 data packs, no problems are encountered. However, if a System/3 Model 6 or Model 10 data pack is used on Model 15 and a nonmultivolume file is attempting to overlay an old System/3 work file and you are using all the original work FILE parameters (LABEL, TRACKS, LOCATION, etc), a terminal error will occur.

When a work file statement is specified on Model 15 systems for 3340s (D1, D2, D3, D4), the verify option is taken from the file statement. Column 34 of the sort header statement is ignored.

Automatic Work File Allocation

If the FILE statement is omitted from the disk sort OCL, disk sort uses the automatic allocation function to allocate from one to a maximum of four scratch areas, each on a different drive. If four scratch areas are not available, disk sort attempts to allocate three, and so on. The scratch areas are allocated from:

1. Available 5445 packs or 3340 areas
2. Available 5444 fixed and removable packs

Note: See chart in the preface of this manual for the disk storage drives applicable to a particular system.

The first choice of the automatic allocation function is the 5445 Disk Storage or the 3340 Direct Access Storage Facility. If a 5445 or 3340 cannot be used, automatic allocation tries to find space on the 5444 disks.

Automatic work file allocation uses scratch areas from the available disk packs. A disk pack is considered available if the pack does not need to be removed while disk sort is executing.

On any given drive, the automatic allocation selects the largest contiguous disk area available. Therefore, a job being executed in another program level or partition may not find enough disk space available. This condition can be minimized to some extent if the disk packs have been fragmented, that is, if a file is intentionally located in the middle of a pack, there will be two contiguous areas instead of one. The degree of pack fragmentation and the location of the fragmented files must be evaluated by each user.

All automatically allocated scratch areas (work file areas) are returned to system use after the disk sort run is completed.

Notes:

1. Disk sort may exhaust the work areas selected by automatic work file allocation. When this occurs, sort issues the message WORK FILE TOO SMALL. To bypass this condition, increase the size of the free areas or specify the work area with a FILE statement.
2. When running a disk sort job using automatic work file allocation, never change packs unless requested to do so by a message.
3. The verify option in column 34 of the header statement applies to any automatically allocated work file.

*Automatic Work File Allocation with Multiple File Input
 (Model 12 and Model 15)*

If more than one input file is assigned to a specific device type (5445 or 3340 or the 5444), then any unit of that type that contains an input file will not be used for automatic work file allocation.

For example, if filename INPUT is assigned to unit D1, INPUT3 to unit R2, and INPUT7 to unit D2, units D1 and D2 are not available for automatic allocation since more than one file is assigned to the 5445.

Note: Unit R2 is available since only one input file was assigned to that device type.

*Directed Automatic Work File Allocation
 (Model 12 and Model 15)*

In order to limit or select the 5444, 5445, 3340, or 3344 drives that are available to disk sort for automatic work file allocation, you can use a SWITCH OCL statement. The SWITCH statement consists of eight characters; one for each of the eight external indicators (U1-U8). The first or leftmost character gives the setting of indicator U1; the second gives the setting of U2; and so on.

File Allocation	Indicator						
	U1&U2	U3	U4	U5	U6	U7	U8
5444 directed	01	X	X	X	R1	R2	F2
5445/3340 directed	11	X	X	D1	D2	D3	D4
X = Not used							

U1 and U2 = 00 No request
 01 5444 directed automatic allocate is requested
 10 No request
 11 5445/3340 directed automatic allocate is requested

U3 and U4 Not used

U5, U6, U7, and U8 When set to a 1, the unit requested is considered available for auto allocation of the work file

When requesting a 5444 directed automatic allocation, U1 must be 0; U2 must be 1; and U6, U7, and/or U8 can be set to 1.

Note: For a 5444 directed automatic allocate, F1 is always considered available for automatic allocation of the work file. If U6, U7, and U8 are all set to 0, then only F1 is considered for automatic allocation of the work file.

When requesting a 5445/3340 directed automatic allocation, U1 and U2 must be set at 1 and at least one of U5, U6, U7, and U8 must be set to 1.

For example, if you want to direct automatic allocation of the work file to D1 and D2, the SWITCH statement would be:

```
// SWITCH 11001100
```

To direct automatic allocation to R1 and F1, the SWITCH statement would be:

```
// SWITCH 01000100
```

The SWITCH statement must be included after the LOAD or CALL statement and before the RUN statement.

Directed automatic allocate does not ensure that space is available on the desired drives; it only directs disk sort requests for work space to these drives.

The size of the work space depends on the number of input records and the size of the record through the intermediate passes. Because space allocation occurs prior to sort generation, estimates of the size of the work files are not known when the space is allocated. Therefore, when automatic allocate is used, it is quite possible that more space will be allocated than is necessary.

If the units indicated on the SWITCH statement for a 5445/3340 directed automatic allocate are not available, automatic allocation takes place on all available 5444 packs.

If the directed automatic work file allocation is used for a job, an appropriate SWITCH OCL statement must be included with the disk sort OCL statements. It is the presence of a SWITCH OCL statement that determines if directed allocation is used. If there is no SWITCH OCL statement, the external indicators are not tested.

When drive 3 (or 4) is specified for a 3344, all main data areas on that drive will be searched for space. For example, if no FILE statement is present and if the SWITCH statement is // SWITCH 11XX0010, disk sort looks for work space only on drive 3. D31 is searched first, then D32, then D33, and finally D34.

If two or more drives are specified, the search for work space alternates. For example, if no FILE statement is present and if the SWITCH statement is // SWITCH 11XX0110, disk sort looks for work space on drive 2 and drive 3. The highest numbered drive, D3 in this example, is searched first. If space for one scratch area (even a single track) is found on D31, then D2 is searched, followed by D31 again, and then D32, D33, and D34 in that order.

Output File

The output file contains the sorted file. The disk file can have sequential organization only. Output of the program is in one of three formats:

- Tags or addrout
- Tag along
- Summary tag along

This page is intentionally left blank.

The file statement name must be OUTPUT. The LABEL keyword can be used to identify the actual file being used.

If the output file exists before the sort (it is old), it cannot have permanent status since you cannot write into a permanent file. However, if the output file is built during the sort, it can have permanent or temporary status.

The output file may be multivolume (online or offline). However, for SORTA jobs, when the input file is processed, the sorted addroot file must be online.

Deferred mounting of the output file is allowed (notice that this does not refer to the system OCL DEFER parameter). This means that the output file does not have to be online prior to running disk sort. The output file specified in the FILE statement may be old or new.

One use of deferred output mounting is to use the same removable unit(s) for both the input and output files, but with the input and output files on different packs. In this case, the shared removable units may not be used for the work file. For example:

```

1   4   8  12  16  20  24  28  32  36  40  44  48  52  56  60
//E
// LOAD $DSORT, F1
// FILE NAME-INPUT, PACK-AAAAAA, UNIT-R1, RETAIN-T
// FILE NAME-WORK, PACK-FIXED1, UNIT-F1, TRACKS-10, RETAIN-S
// FILE NAME-OUTPUT, PACK-BBBBBB, UNIT-R1, TRACKS-10, RETAIN-T
// RUN

```

In this case, pack AAAAAA must be mounted on R1 prior to running disk sort. The mounting of packBBBBB on R1 is deferred until later. R1 may not be used for the work file.

Disk sort defaults to rewinding and unloading tape files.

Variable length blocks are not allowed by disk sort tape I/O; the job ends if they are encountered.

Note: In the output file, the order of records having duplicate control fields is unpredictable. This does not imply that the output is incorrect, however, since the control fields were in order.

7-Track Tape Considerations

If nonstandard and unlabeled tapes are used, extreme care must be taken when creating a file using converter, translator, and parity parameters. Since no labels are given with block and record sizes on this type of tape, identical FILE statements are necessary each time the 7-track tape is used as input for a job. If an error does occur, disk sort cannot recognize the error. No diagnostic is given.

If standard labeled tapes are used, identical FILE statements are still necessary each time the same 7-track tape is used as input for another job. Since there are labels giving block and record sizes, disk sort can diagnose errors if identical FILE statements are not given.

If 7-track tape is used for either input or output, the converter or translator must be used. Refer to the *IBM System/3 Models 8, 10, 12, and 15 Components Reference Manual*, GA21-9236, for a discussion of the converter and the translator.

The translator cannot be used if output is 7-track tape and packed data or summary data specifications are used.

If input is 7-track tape with the translator on, then include, omit, and field specifications with packed data are not allowed.

If SORTA (addrout) sort is specified, output cannot be 7-track tape.

9-Track Tape Considerations

A terminal error occurs if the tape file is ASCII and ASCII-YES is not specified for both the input file and the output file.

If the output for a SORTA (addrout) sort is to be on tape and ASCII translate has been specified, a terminal error is generated and the sort job ends.

For tape output from an addrout sort, the logical record length (RECL-on file statement) must be 18 positions.

Restrictions on the Use of Multifile Tapes (Model 15)

You must adhere to the following restrictions when using multifile tape volumes:

1. All files in the volume must be labeled in the same manner; that is, all must be standard labeled files or all must be unlabeled files.
2. All files in the volume must be recorded in the same density.
3. All files in the volume must be recorded in the same mode (translate, convert, or parity).
4. If the last file on a multifile reel is continued on a subsequent reel, the two reels constitute an aggregate. The restrictions in 1-3 apply to all volumes of the aggregate. In addition, all volumes of the aggregate must be either 7- or 9-track (7-track and 9-track reels may not be mixed in the same aggregate).
5. Standard labeled 7-track tapes, if prepositioned, should be prepositioned to a point just before a HDR1 record. Otherwise, tape data checks or run-away may occur.

Refer to the *IBM System/3 Model 15 System Control Programming Reference Manual*, GC21-5077 and *System/3 Control Programming Concepts and Reference Manual*, GC21-5162, for discussions of tape multifile volumes.

How to Supply OCL Statements for the Model 6

Installations without an Online Data Recorder or a 3741 Data Station: If your installation does not have an online data recorder or a 3741 Data Station, you will use conversational OCL to give the system the information it needs to run the disk sort program.

Appendix A summarizes:

- Which OCL cycle to use (depending on the type of sort you want to do)
- How to respond to the keywords in each cycle

Installations with an Online Data Recorder or 3741 Data Station: If your installation has (1) an online data recorder (IBM 5496 Data Recorder or IBM 129 Data Recorder), (2) a directly attached 3741 Data Station, or (3) a directly attached 3741 Programmable Work Station, you can also supply OCL statements on 96-column punched cards, 80-column punched cards, or 96-byte diskette records.

The *IBM System/3 Model 6 Operation Control Language and Disk Utility Programs Reference Manual*, GC21-7516, explains how to code OCL statements on cards. To supply OCL statements on cards, you must establish the data recorder as the OCL input device. Instructions on how to do this are contained in the *IBM System/3 Model 6 Operator's Guide*, GC21-7501. To supply OCL statements from a diskette, refer to the *IBM System/3 3741 Reference Manual*, GC21-5113.

With or without an online data recorder, OCL statements and error messages will be printed out on whatever output device you specify in your LOG statement at IPL. (See the *IBM System/3 Model 6 Operator's Guide*, GC21-7501, for detailed discussion of the IPL procedure.)

How to Supply OCL Statements for the Model 10

You can supply OCL statements via the device you select as the system reader (the console keyboard, 1442, MFCU1, MFCU2, or 3741). For a discussion of selecting the system reader, see the *IBM System/3 Model 10 Disk System Operator's Guide*, GC21-7508.

To supply OCL statements from the diskette, refer to the *IBM System/3 3741 Reference Manual*, GC21-5113.

How to Supply OCL Statements for the Model 12

You can supply OCL statements via the device you select as the system reader (the console keyboard, 1442, MFCU1, MFCU2, or 3471). For a discussion of selecting the system reader, see the *IBM System/3 Model 12 Operator's Guide*, GC21-5144.

To supply OCL statements from the diskette, refer to the *IBM System/3 3741 Reference Manual*, GC21-5113.

How to Supply OCL Statements for the Model 15

You can supply OCL statements via the device you select as the system reader (1442, 2501, MFCU1, MFCU2, console, MFCM1, MFCM2, or 3741). For a discussion of selecting the system reader, see the *IBM System/3 Model 15 Operator's Guide*, GC21-5075.

TIMING CONSIDERATIONS

The time it takes to sort a file depends on these factors:

1. How much main storage you assign for the program's use. The larger the main storage allocation, the less time it takes to run a sort job.
2. Number of records you want to sort. The greater the number of records, the longer it takes to sort them.
3. Size of the records. The larger the records, the longer it takes to sort them.

4. Number of sequence specifications in the sort program. The greater the number of sequence specifications, the longer it takes to run the sort job.
5. Whether you are using an alternate collating sequence. An alternate collating sequence can increase the time it takes to run the sort job.
6. Where the files are located on the disk. If the disk arm has to make many extra movements, the sort time will increase. This factor is especially significant for jobs with multivolume files.
7. Order of the control fields in the input record.
8. Whether or not automatic work file allocation is being used.
9. Whether you are using the nonverify option. By using this option, data written on the work file will not be verified. The result is a decrease in the time it takes to sort a job.

Each of these factors is discussed in detail in Appendix D

PROGRAMMING CONSIDERATIONS

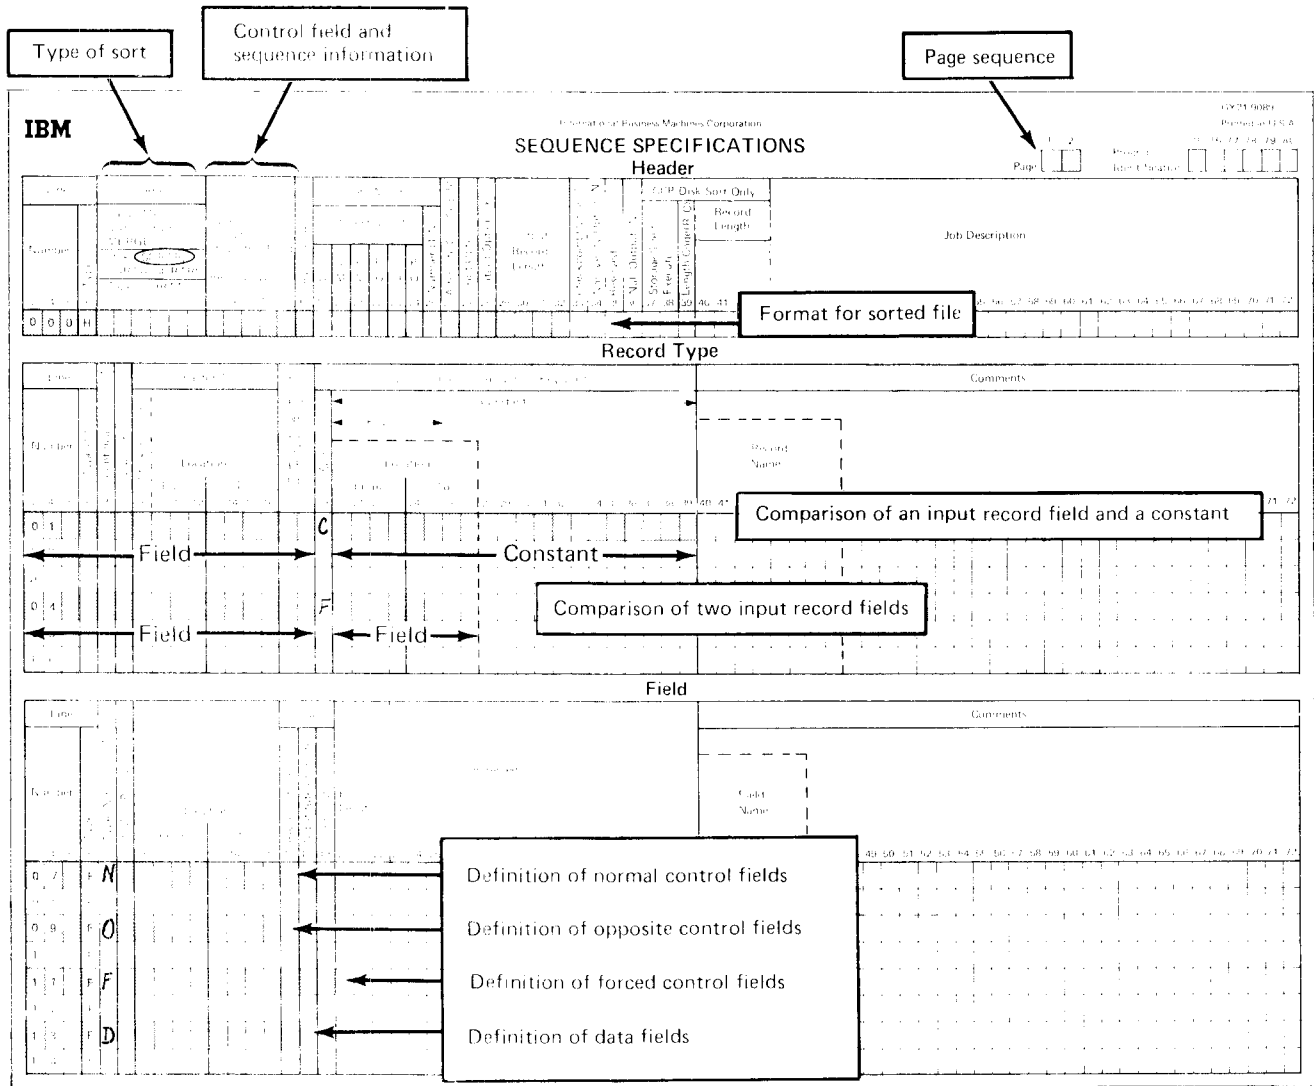
	Addrout Sorts (SORTA)	Tag Along Sorts (SORTR, SORTRS)
Core Storage Requirements	Less than tag along sort	More than addrout sort
Work File Size	Need only be big enough to hold control fields and relative record numbers ¹	Must be big enough to hold entire file
Sorted Output	Relative record numbers only	Sorted output records can contain data only, control fields only, or both data and control fields
Input File	Input file should not be overlaid by output file	Input file can be overlaid by output file ² (the input file can double as the output file)
Job Time	Usually shorter	Usually longer

¹This lets you sort large files without changing packs.

²The space for the output file may be the same as that of the input file. In other words, the output file may overlay the input file. However, good programming practice is to never overlay a file you want to sort, unless you have first made a backup copy of the file. When using the input file as the output file, you must specify the same space (number of records or tracks), location, and label.

SPECIFICATIONS FOR TAG ALONG SORT (SORTR)

The shaded columns are the ones you must consider for a tag along sort job.



SPECIFICATIONS FOR SUMMARY SORT (SORTRS)

The shaded columns are the ones you must consider for a summary sort job.

Type of sort

Control field and sequence information

Page sequence

IBM GX27 9089 1
Printed in U.S.A.

SEQUENCE SPECIFICATIONS

Page 1 / 2 Program Identification 75 76 77 78 79 80

Line		Job	Card Match		Staggered Sort		Output Record Length	Job Description
Number	Type	Location	S	P	M	U	Length	
0 1	F	10 11 12						
0 4	F	13 14 15 16						

Record Type

Line	Type	Factor 1	Rel	Factor 2 (Field, Constant, or Keyword)	Comments
Number	Type	Location	F, O, N, E, L, T, L, F, G, E	From To	
0 1	F	10 11 12	C	Constant	Comparison of an input record field and a constant
0 4	F	13 14 15 16	F	Field	Comparison of two input record fields

Field

Line	Type	Location	Field Name	Comments
Number	Type	From To		
0 7	F	17 18	Normal Control Fields	Definition of normal control fields
0 8	F	19 20	Opposite Control Fields	Definition of opposite control fields
1 1	F	21 22	Forced Control Fields	Definition of forced control fields
1 3	F	23 24	Data Fields	Definition of data fields
1 4	F	25 26	Summary Data Fields	Definition of summary data fields

Chapter 3. Header Specifications

Header specifications tell disk sort:

- What type of sort you want to do (control fields, tags, or summary tag along)
- How you want to format (arrange) the sorted file
- What, if any, system information you want printed (to aid you in error checking)

Remember, use only one header line for each sort job.

COLUMN SUMMARY

Columns	Entries	Explanation
1-2	00-99	Page number
3-5	000	Header line number
6	H	Header line identification
7-12	SORTA	Addroot sort job
	SORTB	Tag along sort job
	SORTBS	Summary tag along sort job
13-17	1-256	Longest control field used in sorting the records
18	A	Records in sorted file to be in ascending order by control fields
	D	Records in sorted file to be in descending order by control fields
19-25	Blank	Not used in disk sort jobs
26	Blank	Use standard System/3 collating sequence in compare operations
	S	Use an alternate System/3 collating sequence in compare operations ALTSEQ statements will define the collating sequence to be used
27	0 or blank	Print: Sequence specification lines Diagnostic messages Program-status messages Action messages
	1	Print: Program-status messages Action messages
	2	Print action messages only
	3	Print nothing
28	Blank	Keep control fields in output records in tag along sort jobs
	X	Drop control fields from output records in tag along sort jobs
29-32	1-4096	Length of output records in tag along sort jobs
33	Blank	Not used in disk sort jobs
34	N	Data written on the work file will not be verified
35-36	Blank	Reserved for system use
37-38	12-48	Storage size to execute the object module in multiples of 2K (Model 15 CCP/Disk Sort [5704-SM7 and 5799-ATH] only)
39	Blank	Reserved for system use except CCP/Disk Sort
	C	Indicates task chain and a record length in columns 40-43 (Model 15 CCP/Disk Sort [5704-SM7] only)
	T	Indicates Release and Task Chain. A record length in columns 40-43 (Model 15 CCP/Disk Sort [5704-SM7] only)
	R	Indicates a record length in columns 40-43 (CCP/Disk Sort only)
40-43	same	Input file record length - right justified with or without leading zeros (CCP/Disk Sort only)
40-72	Any System/3 characters	Comments except CCP/Disk Sort
44-72	Any System/3 characters	Job description for CCP/Disk Sort
73-74	Blank	
75-80	Alphanumeric characters	CCP/Disk Sort object program name

Columns that must be filled in

Columns that must be filled in for a tag along sort

COLUMN DESCRIPTIONS

Columns 1-2 (Page Number) and 3-5 (Line Number)

Page number (columns 1-2) and line number (columns 3-5) form a 5-digit sequence number. As the program reads sequence specifications, it checks the sequence numbers to make sure they are not in descending order. If the numbers are in descending order (for example, if page 02 specifications come before page 01 specifications) and specifications are being printed, the program will print a warning next to the line. After a warning is issued (the letter S) the program continues reading the rest of the specification lines and then halts and waits for further instructions from the operator. The operator can either continue or end the job.

Because page number applies to all lines on a page, columns 1-2 appear only once, in the upper right corner of the page. You number the pages in ascending order.

The line number of the header line is always 000 and is pre-printed on the coding sheet.

Column 6 (Line Type)

Column 6 of the header line contains a preprinted H to identify the line.

Columns 7-12 (Job)

Columns 7-12 tell the program what type of sort job you want to do. SORTA means addrout sort. SORTR means tag along sort. SORTRS means summary tag along sort.

Columns 13-17 [Largest Sum (in Bytes) of Control Field Lengths for Any Record Type]

To calculate this entry:

1. Add together the lengths of the *control* fields (N, O, or F in column 7 of field specifications) for *each* type of input record.
2. Place the *largest* of these totals in columns 13-17. (Your entry must not exceed 256.)

For more information, see *Control Fields* under *Column 7, Field Description Specifications*.

The following specification sheet contains an example of the calculation for the entry in columns 13-17.

SEQUENCE SPECIFICATIONS													Job Description			
Header													Job Description			
Line	Type	Job	Card Match							Output Record Length	Check Input Option	App. Per Y. Option	Reserved	CCP/Disk Sort Only	Record Length	Job Description
Number			Marker	Label	L	S	P	R	S	P	Number	Print Option	Exclude	Record Length		
4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	
1		SortA										3X		128		
2																
3																
4																
5																
6																
7																
8																
9																
10																
11																
12																
13																
14																
15																
16																
17																
18																
19																
20																
21																
22																
23																
24																
25																
26																
27																
28																
29																
30																
31																
32																
33																
34																
35																
36																
37																
38																
39																
40																
41																
42																
43																
44																
45																
46																
47																
48																
49																
50																
51																
52																
53																
54																
55																
56																
57																
58																
59																
60																
61																
62																
63																
64																
65																
66																
67																
68																
69																
70																
71																
72																
73																
74																
75																
76																
77																
78																
79																
80																
81																
82																
83																
84																
85																
86																
87																
88																
89																
90																
91																
92																
93																
94																
95																
96																
97																
98																
99																
100																

Column 18 (Ascending or Descending Sequence)

Control fields control the sequence of records in the sorted output file. Your column 18 entry indicates the sequence in which you want the records sorted:

Col. 18 Entry	Sequence
A	Ascending sequence by control field
D	Descending sequence by control field

Column 26 (Collating Sequence)

Column 26 specifies the collating sequence you want disk sort to use in compare operations. (Compare operations determine whether one character is equal to, greater than, or less than another character.)

Standard Collating Sequence

A blank in column 26 tells the system to use the standard System/3 collating sequence. There are slight variations in the standard collating sequence depending on whether you are using both the zone and digit portions of the characters in your records, the zone portions only, or the digit portions only. Appendix C shows the complete collating sequence for each situation.

Alternate Collating Sequence

An S in column 26 tells the program you want to change the standard collating sequence. To do this you must supply ALTSEQ statements immediately following the header specifications (Appendix C tells you how to code ALTSEQ statements).

Do not use a packed or unpacked factor 1 in an include or omit record type specification (P or U in column 8) if you specify an alternate collating sequence.

Note: Generally, the only users of alternate collating sequences are European firms that want to insert special alphabetic characters (such as the German ä, ö, and ü and the Spanish ñ) into the standard System/3 collating sequence.

Column 27 (Print Option)

The disk sort program can print:

- Sequence specification lines
- Diagnostic messages, for any errors in sequence specifications
- Program-status messages, to identify various stages of the job for you
- Action messages (accompanied by a halt), to identify circumstances requiring attention before you can continue the job

Column 27 indicates which of the preceding information you want the program to print during a job:

Col. 27 Entry	Program Prints
0 or blank	Sequence specifications Diagnostic messages Program-status messages Action messages
1	Program-status messages Action messages
2	Action messages only
3	None of above

Column 28 (Output Option for Tag Along Sorts)

Column 28 applies to tag along sort (SORTR or SORTRS) jobs only. It indicates whether or not you want the program to drop control fields from output records after the records are sorted. A blank in column 28 means keep the control fields; X means drop them.

Considerations for Dropping Control Fields

Control fields are normally dropped if you are using opposite control fields or an alternate collating sequence. In these two cases, the program changes the control information (during the sorting process) in such a way that it will be meaningless to you.

Using Fields as Both Control and Data Fields

If you are using opposite control fields or an alternate collating sequence and you want to keep the control information in a meaningful form in the output records, describe the fields twice: once as control fields and once as data fields. Data fields are not involved in the sorting process and are not changed by the program.

Columns 29-32 (Output Record Length for Tag Along Sorts)

Columns 29-32 apply to tag along sort (SORTR or SORTRS) jobs only. The entry in these columns tells the program the length of records in the final sorted file.

If you do not drop control fields, the length includes both control and data fields. If you drop control fields, the length includes only data fields. In either case, the sum of the control fields and the data fields must not exceed the maximum work record length limits listed in Appendix E.

Calculating Output Record Length When Dropping Control Fields

For *each* type of record, total the lengths of all the data fields you are including in the job. Select the largest total and enter this number in columns 29-32.

Calculating Output Record Length When Not Dropping Control Fields

For *each* type of record, total the lengths of the data fields. Select the largest total, add this total to the number in columns 13-17, and put the sum in columns 29-32. (The sum must not exceed the maximum work record length limits listed in Appendix E.)

Column 34 (Nonverify Option)

Column 34 applies to all disk sort runs except those that contain a work file statement for a 3340 disk file on Model 15 only. When a 3340 file (D1, D2, D3, D4) on Model 15 only is used as the work file, the verify option is taken from the OCL FILE statement and column 34 of the sort header record is ignored. Otherwise, if an N is placed in this column of the disk sort header statement, none of the data written on the work file will be verified. This results in an improvement in performance (see *Appendix D. Timing Considerations*).

Note: If you use the nonverify option when your output file is to overlay the input file, the input file might be destroyed if a terminal error occurs before end of job. Therefore, if the input file cannot be easily recreated, you should have a duplicate copy for backup. The input file can be destroyed under the same circumstances whether the nonverify option is used or not. However, the chance is greater when using the nonverify option.

**Columns 37-38 (Storage Size) – Model 15 CCP/Disk Sort
5704-SM7 and 5799-ATH**

You designate the required storage size, in multiples of 2K, from 12 to 48 for execution of the object module. The default is 12.

**Column 39 (Record Length Indicator) – CCP/Disk Sort
Only**

Models 4, 8, 10, 12, 15A, 15B, and 15C

R This code indicates that a record length exists in columns 40-43.

Model 15 with Program Number 5704-SM7 Only

C Indicates task chain and that a record length exists in columns 40-43.

T Indicates Release and Task Chain operation. A record length exists in columns 40-43. For a discussion of the CCP Release and Task Chain operation program, see *IBM System/3 Communications Control Programmer's Reference Manual, GC21-7579*.

R Indicates that a record length exists in columns 40-43.

Columns 40-43 (Record Length) – CCP/Disk Sort Only

You designate the input file record length, right justified, with or without leading zeros.

Note: If a user comment currently exists, you must fit it into columns 44-72, or remove it.

**Columns 40-72 (Job Description) or Columns 44-72 (Job
Description – CCP/Disk Sort Only)**

This field is for your comments. You can use any System/3 characters you want in these columns. If the program prints specification lines (the column 27 entry is a zero or blank), the comments you include in these columns are printed. Comments have no effect on the program.

Columns 75-80 (Program Identification/Program Name)

Disk Sort

For disk sort, the program identification is only used for user reference. It has no program function.

CCP/Disk Sort

For CCP/Disk Sort, the program name does provide a program function. This field contains the name given to the object program. This is the name that you enter from the terminal to invoke the sort. The entry can be:

1. A valid program name. The first character must be alphabetic, but cannot be a #, \$, or @ character. The remaining characters must be alphameric with no embedded blanks or special characters.
2. Blank. The default is SRTOBJ.

Note: DIR, ALL, and SYSTEM are reserved names and must not be used as program names.

Chapter 4. Record Type Specifications

Record type specifications tell disk sort which of the records in a file you want to sort. Remember, if you want to sort all the records in a file, and they all have the same format, you do not have to fill out record type specifications.

COLUMN SUMMARY

Columns	Entries	Explanation
1-2	00-99	Page number
3-5	01x-06x	Line number (you can leave column 5 blank, or enter any value to keep the specifications in ascending order)
6	I	Include line
	O	Omit line
7	A	<i>And</i> line (these specifications continue the definition of the record described on the previous line)
	O	<i>Or</i> line (these specifications define a different type of record than the one on the previous line)
	Blank	First line of a set of I or O record type lines
	*	Comment line
8	C	Use both zone and digit portions of characters
	Z	Use only zone portion of 1-character field
	D	Use only digit portion of characters
	P	Signed packed decimal data
	U	Signed unpacked decimal data
9-12	1-4096	The input record position in which the factor 1 field begins (blank if field is only one position long)
13-16	1-4096	The input record position in which the factor 1 field ends
17-18	EQ	Factor 1 must equal factor 2
	NE	Factor 1 must not equal factor 2
	LT	Factor 1 must be less than factor 2
	GT	Factor 1 must be greater than factor 2
	LE	Factor 1 must be less than or equal to factor 2
19	GE	Factor 1 must be greater than or equal to factor 2
	C	Factor 2 is a constant
	F	Factor 2 is another field in the same input record
20-23	1-4096	The input record position in which the factor 2 field begins (blank if field is only one position long)
24-27	1-4096	The input record position in which the factor 2 field ends
20-39	Any System/3 characters	The factor 2 constant
40-72	Any System/3 characters	Comments

Columns that must be filled in

COLUMN DESCRIPTIONS

Columns 1-2 (Page Number) and 3-5 (Line Number)

Page number (columns 1-2) and line number (columns 3-5) form a 5-digit sequence number. As the program reads sequence specifications, it checks the sequence numbers to make sure they are not in descending order. If the numbers are in descending order (for example, if page 02 specifications come before page 01 specifications) and specifications are being printed, the program will print a warning (the letter S) next to the line (the S stands for sequence error). After a warning is issued, the program continues reading the rest of the specification lines, then halts and waits for further instructions from the operator. The operator can either continue or end the job.

Because page number applies to all lines on a page, columns 1-2 appear only once, in the upper right corner of the page. Number the pages in ascending order.

Record type line numbers are 01 through 06. The numbers in columns 3 and 4 are preprinted on the coding sheet. You can leave column 5 blank or enter any value to keep your specifications in ascending order. If you have more than six record type lines, use another coding sheet and start at line 01.

Out-of-Sequence Lines

Use column 5 when you want to insert a specification without renumbering the other specifications. For example, to insert a specification line between lines 01010 and 01020, you can number it 01015, code it, and then continue to fill out the rest of the sheet.

Be sure any lines that are out of sequence on your coding sheet are clearly marked. You can do this by writing a note in the margin of the page with an arrow pointing to where the insert belongs (see *Sample Job 3*).

Column 6 (Line Type)

Column 6 identifies the type of record type line. An I in this column stands for either an include or an include-all line; an O stands for an omit line.

Include Line

Include lines identify records you want the program to sort by describing particular record fields.

Include-All Line

Include-all is a special form of include line; it has no record description (columns 7-39 are blank). It tells the program to sort all the records that have not been described by any preceding include or omit line for the job. Records referred to in this manner must have the same field specifications.

Note: Only one include-all line can be used per job. If used, it must be the last record type line for that job.

Summary: Include Lines

If the type of sort job you are running requires you to code record type lines, you must use include or include-all lines to describe the records you want sorted. Records not described in include lines will not be sorted.

Omit Line

Omit lines identify records you do *not* want the program to sort. Omit lines are not required but can be helpful when you have many types of records you want the program to use and just a few you want omitted. Omit lines are normally followed by an include-all line, telling the program to sort all the records that are not described by omit lines.

Sets

There are two types of sets: include sets and omit sets. An include set identifies one or more record types you want to include in your sort job. An omit set identifies one or more record types you want to omit from your sort job. (The records in any record type always have at least one characteristic in common — such as an X in position 5.)

Here are four rules to remember when you are using include and omit sets:

1. All include sets must end with a field description line.
2. Omit sets never have field description lines.
3. Every omit set must be followed by an include set.
4. The last set must be an include set.

There are five types of include sets and three types of omit sets.

INCLUDE SETS¹

Type	Col. 6	Col. 7	Explanation
Include AND lines	H, F, or O		Header line, field description line, or omit line
	I	␣	Record type lines: New record type indicated by blank in column 7.
	I	A	Lines that describe the same record type (as the previous line) have an A in column 7.
	F		Field description line(s)
Include OR lines	H, F, or O		Header line, field description line, or omit line
	I	␣	Record type lines: New record type indicated by blank in column 7.
	I	O	Lines that describe different record types (than the previous line) have an O in column 7.
	F		Field description line(s)
Include AND and OR lines	H, F, or O		Header line, field description line, or omit line
	I	␣	Record type lines: New record type indicated by blank in column 7.
	I	O	This line designates a record type which is different than but has the same field description lines as the record type described in the previous line(s).
	I	A	This line continues the same record type of a previous line or lines. This record type can be continued (IA), or another record type can be started (IO), provided all record types have the same field description lines. Record types with different field lines would have to start a new include set.
F		Field description line(s) for either A or O record types	
Include only one record type (implied include-all)	H		Header line
	F		No record type lines Field description line(s)
Include-all	H, F, or O		Header line, field description line, or omit line
	I		Record type line: Tells program to sort all the records that have not been described by any preceding include and omit lines. Records referred to in this manner must have identical field specifications.
F		Field description line(s)	
<p>¹ Every include set must end with field description lines. An include set can be followed by another include set, an omit set, or // END.</p> <p><i>Note:</i> Records not described in include sets will not be sorted.</p>			

OMIT SETS¹

Type	Col. 6	Col. 7	Explanation
Omit AND lines (one record type)	H or F O O	␣ A	Header line or field description line (last line of include set) New record type indicated by blank in column 7. Lines that describe the same record type (as the previous line) have an A in column 7. } Record Type Lines
Omit OR lines (different record types)	H or F O O	␣ O	Header line or field description line (last line of include set) New record type indicated by blank in column 7. Lines that describe different record types (than the previous line) have an O in column 7. } Record Type Lines
Omit AND and OR lines (different record types)	H or F O O	␣ A O	Header line or field description line (last line of include set) New record type indicated by blank in column 7. Lines that describe the same record type (as the previous line) have an A in column 7. Lines that describe different record types (than the previous lines) have an O in column 7.
¹ There are no field description lines in omit sets. Each omit set must be followed by an include or an include-all set.			

Guide to Using Include and Omit Sets

When to Use Include Sets: If you want to sort only a few records in a file, use an include set for each type of record you want to sort.

When to Use Omit Sets: If you want to sort all but a few records in a file, use omit sets followed by either an include set for each type of record you want to sort or an include-all set.

Mixing Include and Omit Sets: You can mix include and omit sets; but because disk sort processes the sets in the order they are coded, you must be particularly careful when you do this. For example, if you wanted to omit all records with a 2 in position 10 but sort those with a 2 in positions 10 and 15, you would have to specify the include set before the omit set. If you specified the omit set first, all the records you want sorted would be omitted from the job.

Column 7 (Continuation or Comments)

Column 7 indicates the line's relationship to the preceding line.

Col. 7 Entry	Explanation
Blank	This line is the first of a set of include or omit lines. (The type of set is indicated by the column 6 entry: I for include - O for omit.)
A	The line is a continuation of the preceding line. The A stands for AND.
O	The line applies to a different record type than the preceding line, but the control field specifications for both are the same. The O stands for OR.
*	This line is a comment line. Comment lines do not affect the program in any way. Their only purpose is to help you remember what you were doing in a certain section of coding. (Comments are printed only if column 27 of the header line contains a zero or a blank.)

Column 8 (C/Z/D/P/U)

Your column 8 entry tells disk sort how to interpret data in the factor 1 and factor 2 fields during compare operations. When the fields contain alphameric data, a C, Z, or D entry tells disk sort what portions of the characters to use. When the fields contain signed numeric data, a P or U entry tells disk sort whether the data is packed or unpacked.

Col. 8 Entry	Compare Operations	Maximum Field Length ¹
C	Use both zone and digit portions of the characters	256 characters
Z	Use only the zone portion of the character	1 character
D	Use only the digit portion of the characters	16 characters
P ²	Numeric data is packed	8 bytes or 16 digits
U ²	Numeric data is unpacked	16 digits

Alphameric data { C, Z, D

Signed numeric data { P², U²

¹ For both factor 1 and factor 2 fields
² Do not use a packed or unpacked factor 1 in an include or omit record type specification (P or U in column 8) if you specify an alternate collating sequence (S in column 26 of header line).

Significance of the Column 8 Entry

You use record type specifications to tell the disk sort program which records you want to sort. You do this by instructing disk sort to test each record by comparing the data in a specific field against a constant or against the data in another field in the same record. The data you are comparing is the factor 1 field; the data you are comparing it against (a constant or the contents of another field in the same record) is the factor 2 field. The result of the comparisons determines whether or not that record will be sorted. Since the disk sort program sees your data as nothing more than a series of electronic bits, you must tell the program how to interpret the data:

- If the data is alphameric, what part of the characters to use
- If the data is signed numeric, whether it is in packed or unpacked format

Interpreting System/3 Data

Each System/3 EBCDIC character has two parts: a zone portion and a digit portion. Some characters have identical zone portions; some have identical digit portions. No two characters have identical zone and digit portions.

System/3 Character	How It Looks inside the Computer	
	Zone Portion	Digit Portion
*	0101	1100
1	1111	0001
2	1111	0010
3	1111	0011
K	1101	0010
?	0110	1111
P	1101	0111
blank	0100	0000 ¹
0	1111	0000 ¹

¹ Notice that the digit portion of a zero and blank look exactly the same.

If you instruct disk sort to use only the digit portions of characters (by putting a D in column 8), characters with identical digit portions will look alike and compare as equal. Likewise, if you instruct disk sort to use only the zone portion of characters (by putting a Z in column 8) characters with identical zone portions will look alike and compare as equal. Thus your column 8 entry is critical in ensuring that your compare operations produce the results you intend.

Suppose, for example, that you want only those records with a 2 in column 15 and a 2 in column 50. If you put a D in column 8, you will get the records you want: the ones with a 2 in column 15 and a 2 in column 50. But you will also get a lot of records you do not want (several characters have the same digit portion as a 2). To get *only* the records with a 2 in column 15 and a 2 in column 50 you would have to put a C in column 8. The C tells disk sort to use both the zone and digit portions of characters in its compare operations, and no other System/3 character has the same zone and digit portion as a 2.

Note: The leftmost bit of each byte of ASCII data that is not converted to EBCDIC is set to zeros.

Alphameric Data

When the factor 1 and factor 2 fields contain alphameric data, the column 8 entry must specify one of the following:

- The zone and digit portions of the characters (C entry)
- Only the zone portion of the character (Z entry)
- Only the digit portion of the character (D entry)

Numeric Data

When the factor 1 and factor 2 fields contain numeric data, the column 8 entry must answer two questions:

1. Are the numbers signed or unsigned?
2. Are the numbers packed or unpacked?

Format of Numeric Data		Col. 8 Entry	Definition	Considerations ¹
Unsigned		D	Number does not have a sign.	Absolute values of the numbers are used. If a number has a sign, it is ignored. For example, -3 would be considered equal to +3.
Signed	Packed	P	Number always carries a sign. When the number is placed in core storage it has a digit portion only.	The sign controls the comparison. For example, -3 is less than 0 and +5 is more than -6.
	Unpacked	U	Number always carries a sign. When the number is placed in core storage, it has both a zone and digit portion	The sign controls the comparison. For example, -3 is less than 0 and +5 is more than -6.

¹ Before comparing numeric data, disk sort converts any leading blanks to zeros.

Signed Numbers: Signed numbers can be either positive or negative. The sign of a number is indicated by a 4-bit binary code.

Packed and Unpacked Numbers: An unpacked digit takes up eight bits; a packed digit takes up four bits.

Sign	Binary Code
+	1111 ¹
	1010
	1100
	1110
-	1101
	1011

¹ Standard form. Disk sort accepts all four forms of the plus sign. Before sorting the file, however, the program converts all plus signs to the standard form. If you print a core dump, the plus sign will always be expressed as a hex F.


Packed Digits	
Digit	Binary Form
0	0000
1	0001
2	0010
3	0011
4	0100
5	0101
6	0110
7	0111
8	1000
9	1001

Digit Portion Only

Unpacked Digits		
Digit	Binary Form	
0	1111	0000
1	1111	0001
2	1111	0010
3	1111	0011
4	1111	0100
5	1111	0101
6	1111	0110
7	1111	0111
8	1111	1000
9	1111	1001

Zone Portion Digit Portion

In unpacked numbers, the 4-bit sign code replaces the zone portion of the last digit in the number. In packed numbers the sign code takes up the last four bits of the number.



Columns 9-16 (Factor 1 Location)

Factor 1 fields identify your records. (If all your inventory records contain an I in column 2, for example, column 2 is a factor 1 field.) Disk sort identifies records you want sorted by comparing factor 1 fields against constants or against other fields in the same record. Columns 9-16 identify the locations of the factor 1 fields in the records. If there is more than one factor 1 field for the records you are describing, you must:

- Describe each field in a separate record type line.
- Put an A in column 7 of every line (except the first) to tell disk sort that all the lines apply to the same record type.

Columns 9-12 (from) identify where the factor 1 field begins in the record. Columns 13-16 (to) identify where the field ends.

Length of Factor 1 Fields

A factor 1 field can contain anywhere from 1 to 256 characters. No factor 1 field, however, can be longer than the length of the records you are working with. (For example, when you are working with 96-column records, the longest possible factor 1 field you can have is a 96-character field.) The length of factor 1 fields is also controlled by the column 8 entry.

Col. 8 Entry	Maximum Factor 1 Field Length
C	256 characters ¹
Z	1 character
F	16 characters
P	8 characters ²
L	16 characters
¹ When factor 2 is a constant, the length of the factor 1 field must not exceed 20 characters (see Columns 20-39 for more information). Because the factor 1 field is packed, the field can actually represent 15 decimal digits and a sign.	

Coding Rules

Entries must be right-justified: the from entry must end in column 12; the to entry must end in column 16.

To describe factor 1 fields that are only one character long, leave columns 9-12 (from) blank and enter the number of the record position that contains the character in columns 13-16 (to).

For example, here is the entry you would make to describe the position of a factor 1 field which consists of an I in column 2.

Columns 17-18 (Relation)

The program identifies records you want to sort by comparing the factor 1 field (columns 9-16) against either a constant (columns 20-39) or another field (columns 20-27) in the same record. The constant or other field is called factor 2. Columns 17-18 tell the program what the results of the comparison must be. (If alternate collating sequence is used, both factor 1 and factor 2 are modified before the comparison is made.)

Col. 17-18 Entry	Meaning
EQ ¹	Factor 1 must equal factor 2.
NE ¹	Factor 1 must not equal factor 2.
LT	Factor 1 must be less than factor 2.
GT	Factor 1 must be greater than factor 2.
LE	Factor 1 must be less than or equal to factor 2.
GE	Factor 1 must be greater than or equal to factor 2.
¹ If you want the program to compare zone portions of characters (Z in column 8), EQ and NE are the only entries you can use.	

Column 19 (Field or Constant)

The program identifies records you want to sort by comparing the factor 1 field (columns 9-16) with a constant, or another field in the same record. The constant or other field is called factor 2. Column 19 tells the program whether factor 2 is a constant or another field. C in column 19 means factor 2 is a constant; F means factor 2 is a field.

When you put a C in column 19, you use columns 20-39 for the constant. When you put an F in column 19, you use columns 20-27 to identify the location of the factor 2 field in the records.

Note: The K shown in column 19 on the sequence specification sheet is not used for System/3 disk sort.

Columns 20-27 (Factor 2 Field)

The factor 2 field must be the same length as the factor 1 field. It also must be in the same record as the factor 1 field.

Columns 20-27 are used to record the location of the factor 2 field. Columns 20-23 (from) identify the starting position of the field; columns 24-27 (to) identify where the field ends.

Note: The keyword parameter shown for columns 20-25 on the sequence specification sheet is not used for System/3 disk sort.

Coding Rules

Entries must be right-justified: the from entry must end in column 23; the to entry must end in column 27.

To describe fields that are only one character long, leave columns 20-23 (from) blank, and enter the number of the record position that contains the character in columns 24-27 (to).

Columns 20-39 (Factor 2 Constant)

When factor 2 is a constant, you use columns 20-39 to write in the constant you want to use. The constant can be any arrangement of System/3 characters.

The constant must be the same length as the factor 1 field. For example, if you have a 4-position factor 1 field, your constant field must take up four positions. If your constant is the number 6, you would put the 6 in column 23, and either leave columns 20, 21, and 22 blank or fill them with zeros.

		Record Type																					
		Factor 2 (Field, Constant, or Keyword)																					
		Keyword												Constant									
Column	From	To	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	
19																							

or

		Record Type																					
		Factor 2 (Field, Constant, or Keyword)																					
		Keyword												Constant									
Column	From	To	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	
19																							

If the factor 1 field contains a packed number, the length of the constant (including the sign) must be twice the length of the factor 1 field. The reason for this is that factor 1 data is in packed form, and the constant you are writing in is in unpacked form.

Alphameric Constants (Column 8 Entry is C, Z, or D)

The constant must be the same length as the factor 1 field and must always begin in column 20.

Numeric Constants (Column 8 Entry is P, U, or D)

Format: Numeric constants must be right-justified within the field length specified in factor 1 (within twice the field length if factor 1 is a packed number). For example, assume that factor 1 defines a 6-position field in the input record, and that factor 2 is the numeric constant 123. To right justify the constant within six positions, you would have to put the constant in columns 23, 24, and 25. Leading zeros are not required. To disk sort, blanks and zeros look the same. In the example above, columns 20-25 could contain either 000123 or bbb123 (with b representing a blank).

Signed Constants: If factor 1 is a packed number, the last character in the constant must be its sign (+ or -). If factor 1 is an unpacked number and the constant is a negative number, the last digit in the constant must be a character that indicates both the numeric value of the last digit and the negative sign for the entire constant.

Coding Negative Unpacked Constants			
If Last Digit in Constant is	Character that You Code	How the Number Looks inside the Computer	
		Zone Portion ¹	Digit Portion ²
0	- (minus code) ³	0110	0000
1	J	1101	0001
2	K	1101	0010
3	L	1101	0011
4	M	1101	0100
5	N	1101	0101
6	O	1101	0110
7	P	1101	0111
8	Q	1101	1000
9	R	1101	1001

¹The zone portion indicates the negative sign of the entire number.
²The digit portion indicates the numeric value of the last digit in the number.
³If your system input device is a system card device, you can punch a zero and minus in the same card column.

For example, here is the entry you would make to sort records which have a packed negative 1 (-1) in positions 1 and 2, an unpacked negative 24 (-24) in positions 5-8, and an unpacked negative 10 (-10) in positions 11-16:

Record Type																
Line		Factor 1				Ret		Factor 2 (Shifts, Constants, Systems)								
Number		Type (I/O)	Character Range	Location	From	To	Location	From	To	Location	From	To	Location	From	To	
3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	
01		I	P	1	2	EQC		1-								
02		I	A	5	8	EQC		2M								PACKED -1
03		I	O	11	16	EQC		1-								UNPACKED -24
04																UNPACKED -10

Columns 40-72 (Comments)

Columns 40-72 are for your comments. If you have instructed disk sort to print sequence specifications (column 27 of the header line is either blank or contains a zero) the comments will be printed along with your sequence specifications. The comments have no effect on the program's operation.

Many programmers like to use columns 40-47 to write the names of the records described in the record type specifications. This is why columns 40-47 are enclosed in dotted lines on the coding sheet.

Field description specifications tell disk sort how to arrange (format) records in the output file.

COLUMN SUMMARY

Columns that must be filled in for all sort jobs

Columns that must be filled in when forced control fields are used

Columns	Contents	Explanation
1-2	00-99	Page number
3-5	07x-14x	Line number (you can leave column 5 blank, or enter any value to keep the specifications in ascending order).
6	F	Field specification line
7	N	Normal control field
	O	Opposite control field
	F	Forced control field
	D	Data field
	S	Summary data field
	*	Comment line
8	P	Signed packed decimal data
	U	Signed unpacked decimal data
	C	Use both zone and digit portions of characters in the field
	Z	Use only zone portion of 1-character field
	D	Use only digit portion of characters in the field
	V	Force a data character into the data field
9-12	1-4096	Starting position of field in the record (blank if field is one character long)
13-16	1-4096	End position of field in the record
17	Any System/3 character	Forced control fields only (the character you want the program to change)
18	Any System/3 character	Forced control fields only (the character you want to substitute)
19	Blank	Forced control field line is not a continuation of the preceding line
	Any character other than blank	Forced control field line is a continuation of the preceding line
20-22	1-16	Summary tag along sort only (overflow field length entry)
23-39	Not used	Not used
40-72	Any System/3 characters	Comments

COLUMN DESCRIPTION

Columns 1-2 (Page Number) and 3-5 (Line Number)

Page number (columns 1-2) and line number (columns 3-5) form a 5-digit sequence number. As the program reads sequence specifications, it checks the sequence numbers to make sure they are not in descending order. If the numbers are in descending order (for example, if page 02 specifications come before page 01 specifications) and specifications are being printed, the program will print a warning next to the line. After a warning is issued (the letter S), the program reads the rest of the specification lines and then halts and waits for further instructions from the operator. The operator can either continue or end the job.

Because page number applies to all lines on the page, columns 1-2 appear only once, in the upper right corner of the page. Number the pages in ascending order.

Here is the recommended procedure for field description line numbers. Field line numbers are 07 through 14. The numbers in columns 3 and 4 are preprinted on the coding sheet. You can leave column 5 blank, or enter any value to keep your specifications in ascending order. If you have more than eight field lines, use another coding sheet and start at line 07.

Use column 5 when you want to insert a specification line without renumbering the other lines. For example, to insert a specification line between lines 01070 and 01080, you can number it 01075, code it, and then continue to fill out the rest of the sheet.

Be sure any lines that are out of sequence are clearly marked. You can do this by writing a note in the margin of the page with an arrow pointing to where the insert belongs (see *Sample Job 3*).

Column 6 (Line Type)

Column 6 contains a preprinted F, identifying the line as a field line. For addroot sorts (SORTA), field lines describe control fields the program uses to sort record addresses. For tag along sorts (SORTR) and summary tag along sorts (SORTRS), field lines describe the fields that the program uses to create the records in the sorted output file. The fields can be either control fields (used to sort the records) or actual data fields. In addition, for summary tag along sorts, field lines describe the fields that the program summarizes (adds together).

Column 7 (Field Type or Comments)

Your column 7 entry tells disk sort whether you are describing a control field, data field, or comment line. If you are describing a control field, the column 7 entry indicates how the field is to be used. See *Columns 7-8* to find the ways you can combine this entry with the column 8 entry.

Col. 7 Entry	Tells Program
D	This is a data field ¹
*	This is a comment line.
N	This is a normal control field. Sort this field so that the data from the field is in the sequence specified in column 18 of the header line.
O	This is an opposite control field. Sort this field so that the data from the field is in the sequence opposite that specified in column 18 of the header line.
F	This is a forced control field. Change the control field according to the entries in columns 9-19.
S	This is a summary data field ²
¹ Use this entry for tag along (SORTR, SORTRS) sorts only. (If you use a D entry during an addroot sort, the line will be treated like a comment line.)	
² Use this entry for summary tag along (SORTRS) only. (If you use an S entry during an addroot sort, the line will be treated like a comment line. If you use an S entry during a tag along sort, the fields will be treated as normal data fields.)	

Data Fields (D in Column 7)

Data fields apply to tag along (SORTR, SORTRS) sort jobs only. They are fields you want the program to include in the sorted records, but which you do not want the program to use in sorting the records. Within each include set (include and field description lines), control field lines must be placed before data field lines.

When your file has more than one type of record:

- The number of data fields does not have to be the same for all record types.
- The total lengths of all the data fields do not have to be the same for all record types. Disk sort places blanks to the right of shorter data fields so that all total lengths are equal.

Comment Lines (in Column 7)*

Comment lines help document the program. They do not affect the program's operation. You can code comment lines anywhere in the sequence specifications; however, comments will be printed only if column 27 of the header line contains a zero or blank.

Control Fields (N, O, or F in Column 7)

There is no limit to the number of control fields allowed, however, the total length of all the control fields must be from 1 to 256 bytes long.

When your file has more than one type of record:

- The number of control fields does not have to be the same for all record types.
- The total *lengths* of the control fields do not have to be the same for all record types.

When records with duplicate control fields (records with all control fields equal) are sorted, their order as output is unpredictable.

Normal and Opposite Control Fields (N or O in Column 7): These control fields are fields the program uses to sort records or record addresses. They are fields in your input records. However, you can define 1-character control fields that are not in the records by using an unconditional force. (See *Forced Control Fields*.)

Forced Control Fields (F in Column 7): There are three types of forced control fields:

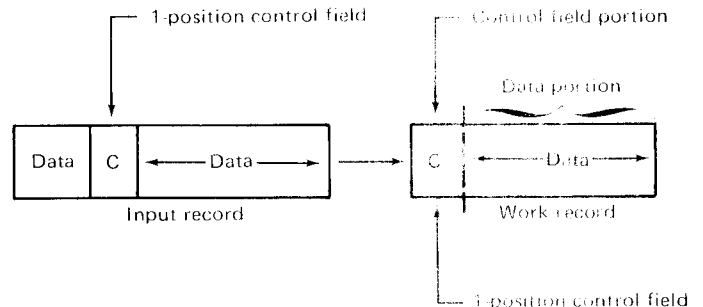
- Conditional
- Force all
- Unconditional

Forced control fields affect the work and output records only. (It is important to remember that disk sort never changes your input records.)

See *Column 17, Column 18* and *Column 19* for information on how to fill out the field description specifications when you are using forced control fields.

A conditional force occurs *only* if a control field in the input record contains a particular entry. Suppose, for example, that you want to sort a file of records, each of which has a 1-position control field. If the character in the control field is an X, you want to replace it with an A before you sort the records. To do this, you would use a conditional force. Your conditional force field would tell disk sort:

- Build a work record from the input record:



- If the control field contains an X, change it to an A:




Force-all is a special type of conditional force. Force-all occurs only when a control field in an input record *does not* contain a particular entry. Suppose, for example, that you want to sort a file of records, each of which has a 1-position control field. If the character in the control field is *not* a C, you want to put an X in the field before you sort the records. To do this, you would use a force-all line to force the X into the control field.

A force-all line follows a series of conditional force lines. For example, you may want to tell disk sort:

- If the control field contains a C, replace it with a 1.
- If the control field contains an F, replace it with a 2.
- If the control field contains a \$, replace it with a 3.
- If the control field does *not* contain a C, an F, or a \$, put an X in the control field. (You are *forcing out all* other possible entries by using a force-all line.)

An unconditional force does not depend on entries in the input records. If, for example, you want to put a dollar sign (\$) in the first position of every output record, you would use an unconditional force. Your unconditional forced control field would tell disk sort to put a dollar sign in the first available control field position of the work record.



Any other control fields will go after the dollar sign. Assume the input records are in this format:

Data	Control field A	Data	Control field B	Data
------	-----------------	------	-----------------	------

The work record will then look like this:

\$	Control field A	Control field B	Data
----	-----------------	-----------------	------

See *Column 19* for actual coding examples of forced control fields.

Summary Data Fields (S in Column 7)

An S in column 7 of the field description specifications defines a summary data field. You can define summary fields for all three types of sort jobs, but the fields will be summarized (added together) only in summary sorts. In tag along sorts (SORTR), summary fields will be treated as normal data fields. In addrout sorts (SORTA), summary fields will be treated as comments.

In a summary sort, the summary data fields in the work and output records for individual record types must always be in the same position. The fields do *not* have to be in the same position in the input records. (See messages SD278 and SD280.)

No more than 24 fields can be summarized for each record type.

The first include set that contains summary specifications defines the summary format for all included records. It is recommended that all include sets contain summary specifications. If an include set does not contain summary specifications, the data specification should align the data for summarization.

If no summary specifications are given for a SORTRS type sort, the output file consists of records with unlike control fields as specified in the field statement. Disk sort deletes all records with common control fields. Only one copy of each is retained.

Column 8 (C/Z/D/P/U/V)

Your column 8 entry indicates what portion of the input record's characters you want disk sort to use in building and sorting the work records. The column 8 entry is critical in assuring that the sort produces the results you intend. See *Columns 7-8* to find the ways you can combine this entry with the column 7 entry.

Col. 8 Entry	Character Portion Used	Maximum Field Length
C	Use both zone and digit portions of the characters	256 characters
Z	Use only the zone portion of the character	1 character
D	Use only the digit portion of the characters	16 characters
P	The characters are signed, packed decimal numbers ¹	8 bytes or 15 digits and sign
U	The characters are signed, unpacked decimal numbers ¹	16 digits
V	Force a data character into the data field	1 character

¹ Note that -3 is less than 0 and +5 is greater than -6.

Suppose, for example, you have a 1-character control field in your input records which can be either an *, 1, 2, or 3. The zone and digit portion of each character is:

Character	Zone	Digit
*	0101	1100
1	1111	0001
2	1111	0010
3	1111	0011

If you want the records sorted into ascending order using the digit portion of the control field characters (by putting a D in column 8), they will be in this order:

- 1
- 2
- 3
- *

If you want the records sorted into ascending order using both the zone and digit portions (by putting a C in column 8), they will be in this order:

*
1
2
3

Suppose you placed a Z in column 8 and wanted the records sorted into ascending order. You can then be sure that the records with an * control field will precede the records with a 1, 2, or 3 control field. Since 1, 2, and 3 have identical zone portions, records with any of these numbers as a control field will not be in any special order after the sort.

If you want to force characters into your data field, place a V in column 8 and specify the character to be forced in column 18. That character will be placed in the first available data field position of the work record.

As you can see, your column 8 entry can drastically affect your sorted file.

Do not confuse this column 8 entry with the column 8 entry on the record type specifications. Column 8 of the field type specifications tells disk sort what portion of a character to use to sort the records. The column 8 entry on the record type specifications helps select which records you want to be sorted.

Opposite Control Fields

If you want to sort records so that some control fields are in ascending order and other control fields are in descending order, use opposite control fields. An opposite control field is sorted in ascending order (if you specify descending order on the header line), or in descending order (if you specify ascending order on the header line).

If your file contains different record types, all of which have an opposite control field in the same record position, your column 8 entries for these control fields must be one of the following:

- All Ds
- All Cs
- All Zs
- Any combination of Cs and Zs

With any other combination of entries (for example, Ds and Cs), you will not be able to predict the results of the sort.

When you use opposite control fields, disk sort changes them in building the work record. Therefore, you usually drop this meaningless control field information (by coding an X in column 28 on the header line) for tag along or summary sorts. If the opposite control fields are all Ds, you do not need to drop the control field. If you wish to retain the original control field data in the output record, repeat the information as a data field.

Packed or Unpacked Control Fields (Normal or Opposite)

If you specify packed or unpacked control fields, disk sort changes the control fields while building the work record. Therefore, you must drop the control field information by coding an X in column 28 on the header line. If you wish to retain the original control field data in the output record, repeat the information as a data field.

Forced Control Fields

Here are a few rules to remember when you are using forced control fields:

- You can have only one character in your forced control field.
- You can indicate either a conditional or an unconditional force.
- A force-all line must be preceded by a conditional force line.
- You define a forced control field by placing an F in column 7 of the field specifications.

See *Columns 9-16*, *Column 17*, *Column 18*, and *Column 19* for further information on how to complete the field specifications.

Using Control Fields to Sequence Information in the Sorted Records

The order in which you describe control fields in the field specification lines determines the sequence of the records (tag along sort) or the record addresses (addrout sort) in the sorted file.

Suppose each record in your file that is sorted in ascending order (A in column 18 of the header line) has a normal control field in positions 1-2 and an opposite control field in positions 5-7. Each record represents one customer's order for a separate item. The part number is in position 1-2; the number of parts ordered is in positions 5-7. Your unsorted file might look like this:

		Input record position						
		1	2	3	4	5	6	7
0		1	7			1	2	6
1		3	4			5	1	2
2		1	7			1	2	1
3		1	5				9	6
4		3	4				1	8
5		1	7				2	0
6		3	4			2	5	0

Input record number Part number Number ordered

You can use the first control field to sort the records in ascending order according to the part number. You can then use the second control field to sort the number of parts ordered in descending order within each group of parts. Therefore, you want your sorted file to look like this:

		Output record position				
		1	2	3	4	5
0		1	5		9	6
1		1	7	1	2	6
2		1	7	1	2	1
3		1	7		2	0
4		3	4	5	1	2
5		3	4	2	5	0
6		3	4		1	8

Output record number Part number Number ordered

Include these field specifications to do the preceding sort:

IBM SEQUENCE SPECIFICATIONS Header

Line	Number	Field	Location	Length	Code
0 1					
0 2					
0 3					
0 4					
0 5					
0 6					
0 7					
0 8					
0 9					
1 0					

Handwritten in table: SORTR, 5A, 0X, 5

Record Type

Line	Number	Field	Location	Length	Code
0 1					
0 2					
0 3					
0 4					
0 5					
0 6					
0 7					
0 8					
0 9					
1 0					

Handwritten in table: *

Field

Line	Number	Field	Location	Length	Code
0 1					
0 2					
0 3					
0 4					
0 5					
0 6					
0 7					
0 8					
0 9					
1 0					

Handwritten in table: WC 1 2, OC 5 7, DC 1 2, DC 5 7

Columns 7-8

Handwritten notes: ALL RECORDS INCLUDED IN SORT
THIS CONTROL FIELD IS DROPPED
THIS CONTROL FIELD IS DROPPED
CONTROL FIELD DESCRIBED AS DATA
CONTROL FIELD DESCRIBED AS DATA

Notice that the information occupies different positions in the input and output records. The record positions change when the control fields and data fields are moved from the input to the work to the output files. Since an opposite control field is specified, all the control fields in the input record are dropped (columns 1-2 and 5-7). To save the control fields, specify them as data (see example).

Here are all the possible combinations for columns 7 and 8:

Col. 7	Col. 8	Maximum Field Length
N or O	C	256
	Z	1
	D	16
	P	8
	U	16
F	C ¹	1
	Z ¹	1
	D ¹	1
D	C	256
	Z	1
	D	16
	P	8
	U	16
	V	1
S	C	256
	Z	1
	D	16
	P	8
	U	16
	V	1
	Does not apply (comment line)	

¹ For an unconditional force and a force-all line, column 8 must contain a C.

See *Column 7* and *Column 8* for detailed discussions of each entry.

Defining a Conditional Force Character

1. Fill in columns 1-6 as you would for any control field.
2. Put an F in column 7.
3. Define the position of the control field in the input record in columns 13-16.
4. Enter the character you want to replace in column 17.
5. Enter the character you want to replace it with in column 18 (you can use any System/3 character in column 18).

Defining a Force-All Character

1. Fill in columns 1-6 as you would for any control field.
2. Put an F in column 7.
3. Put the character which replaces the control field in column 18.
4. Put any System/3 character in column 19. (The character in column 19 tells disk sort that the line is a continuation of the preceding line.)
5. Leave columns 9-17 blank.

If you do *not* place a force-all line after conditional force lines and disk sort does *not* find the specified characters in the control field of the input record, disk sort:

- Replaces the control field character with hex FF (if you specified ascending sequence in the header line).
- Replaces the control field character with hex 00 (if you specified descending sequence in the header line).

Defining an Unconditional Force Character

1. Fill in columns 1-6 as you would for any control field.
2. Put an F in column 7.
3. Put the character you are forcing in column 18.
4. Leave columns 9-17 blank.

Sample Job 4 illustrates the use of conditionally forced characters to group record types.

Specifying an Overflow Indicator Field

1. Fill in columns 1-6 as you would for any control field.
2. Put an S in column 7.
3. Put a V in column 8.
4. Enter a character in column 17 if you do not wish to use an * as a replacement. If overflow occurs in any summary data field in the record, the initial value of the overflow field will be replaced by the character in column 17.
5. Enter a character in column 18. If none of the summary data fields overflow, the overflow field will contain the character specified in column 18.

Column 19 (Specifying Replacements for More Than One Character of a Forced Control Field)

If a control field in the input record can contain any one of several characters and you want to specify forced character replacements for more than one character, use a separate line to define each possible character and the forced character you want to replace it with.

Note that column 19 is assumed blank for the first field statement of an include set.

Coding Rules

1. Define the control field in columns 13-16.
2. Enter the first character you want to replace in column 17.
3. Enter the character you want to replace it with in column 18.
4. Enter the next character you want to replace in column 17 of the next card.
5. Enter the character you want to replace it with in column 18.
6. Enter any character (except blank) in column 19. This tells disk sort the line refers to the same control field in the work record as the preceding line.

Repeat steps 4, 5, and 6 for any other characters to be replaced in that control field.

Note that you can use any System/3 character you want in column 18.

Examples of Using Forced Control Fields


You can use forced control fields to change (add to or replace) a character in a 1-position control field. The following examples represent these types of forced control characters:

- Unconditional force
- Conditional force using only forced control fields

- Conditional force using normal or opposite and forced control fields
- Force-all

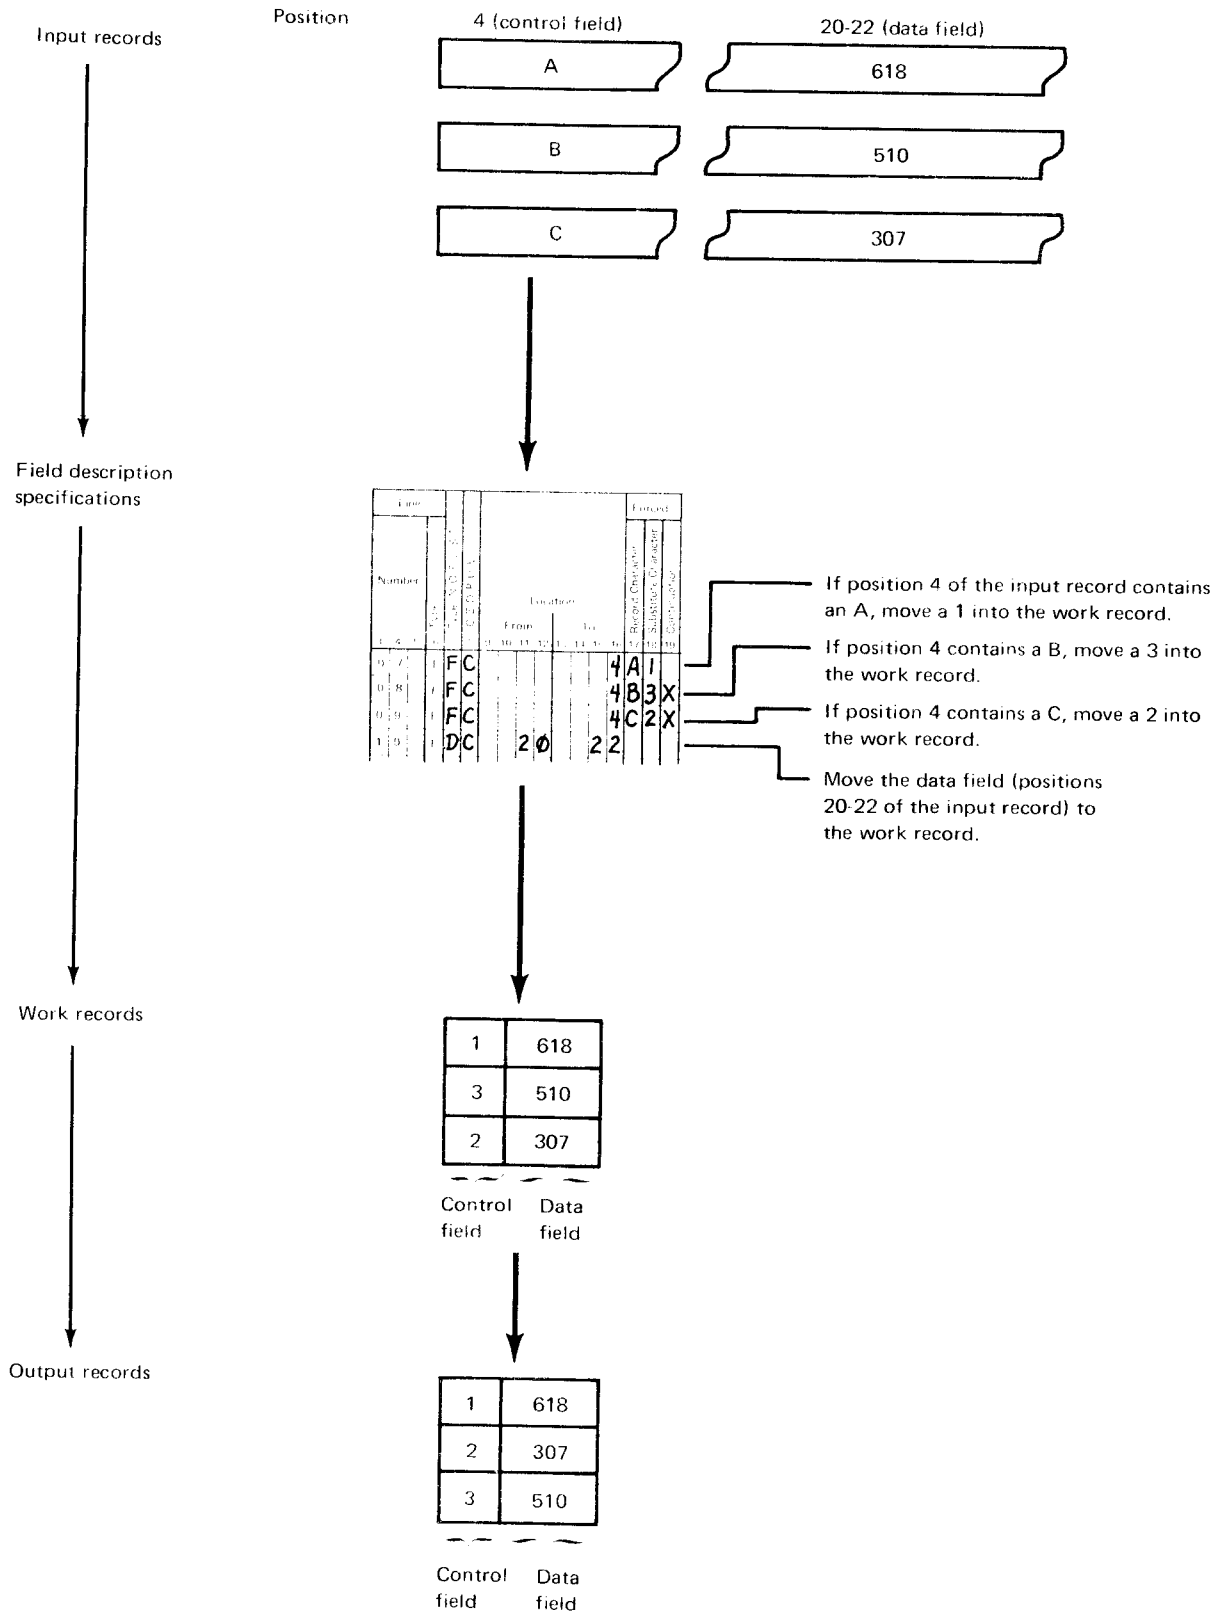
Unconditionally Forced Character

This example illustrates how you can unconditionally place a control field into the work and output records:



Conditional Force Using Only Forced Control Fields

This example illustrates how control field characters change according to the conditions you specify.



This example also shows how you use column 19 to specify conditional replacements for more than one character. Because continuation lines are used, only one position in the work and output records is defined by the first three lines. If column 19 is left blank, each line would define a new position in the work and output records.

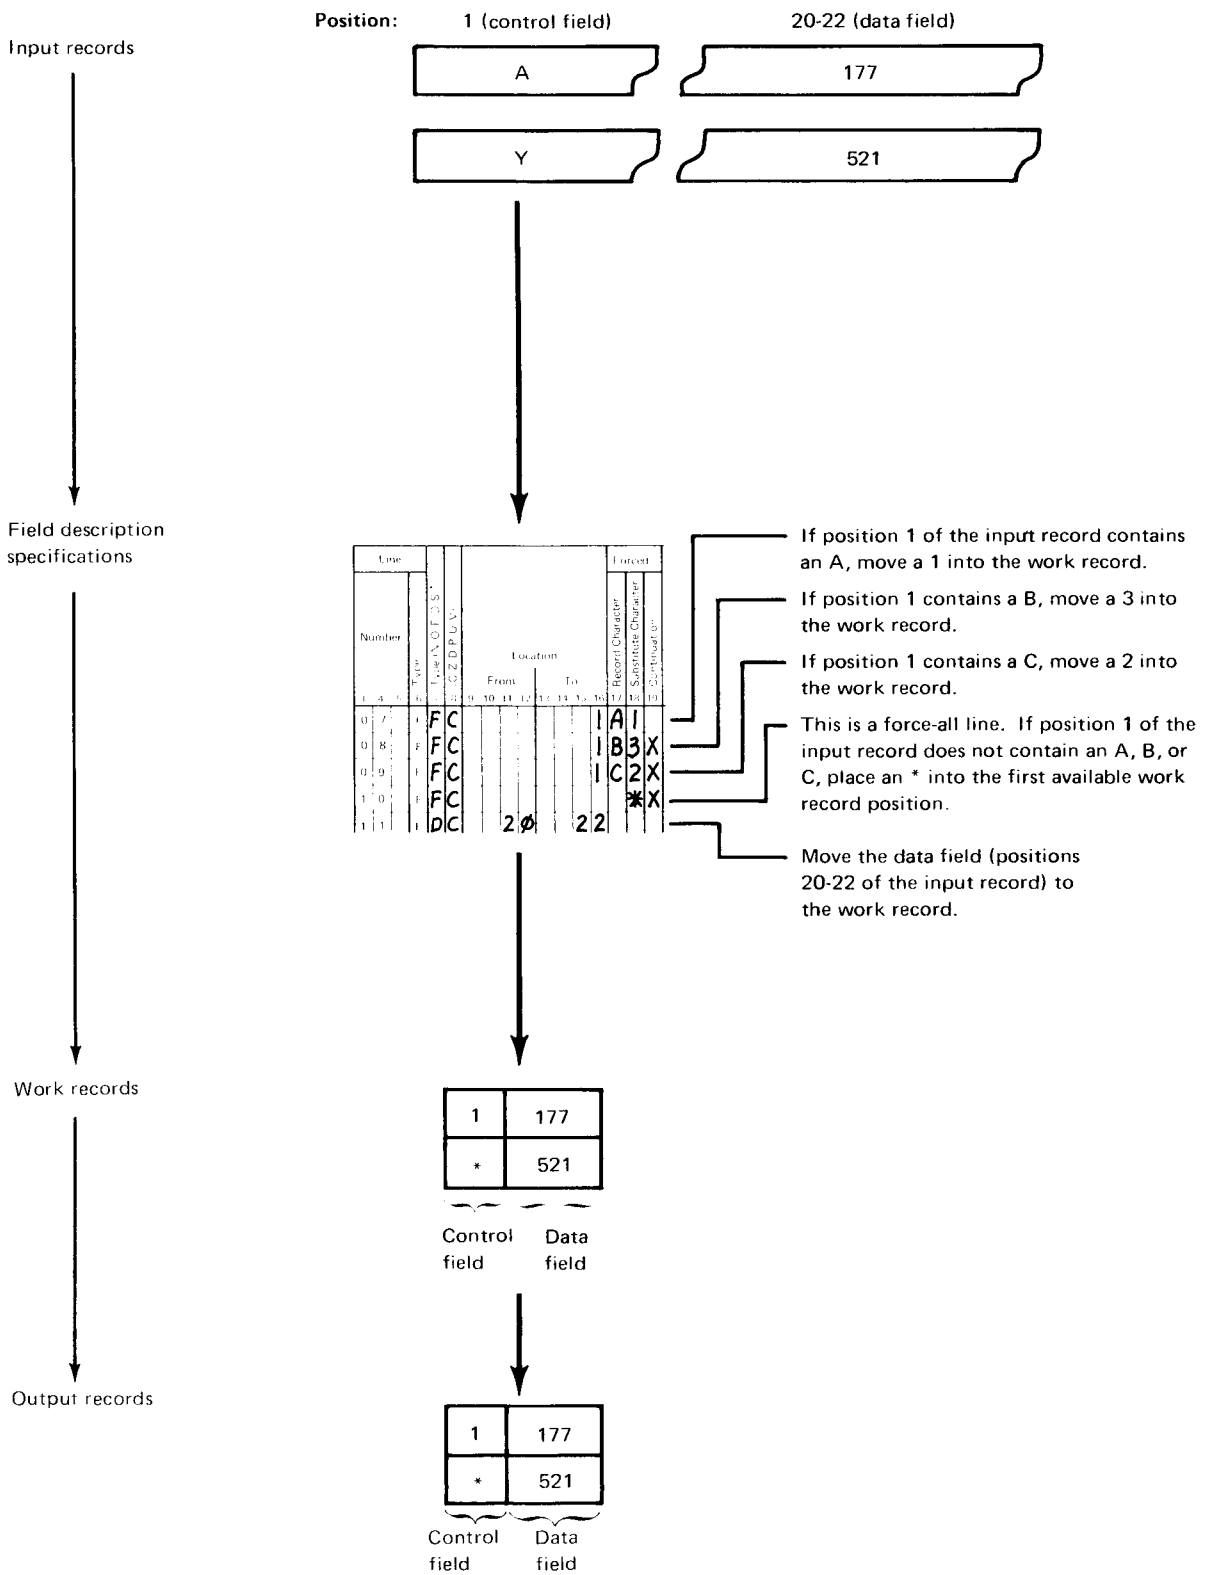
If you are sorting records into ascending order (A in column 18 of the header line), disk sort places hex FF into the work record before you force any characters. If you are sorting records into descending order (D in column 18 of the header line), disk sort places hex 00 into the work record before you force any characters. Therefore, if disk sort does not find an A, B, or C in position 4 of the input record, the hex 00 or FF is left unchanged in the work record.

Conditional Force Using Normal or Opposite and Forced Control Fields

This example is similar to the previous one since the control field in the input record causes changes in the work and output records. In this case, however, disk sort first moves the control field to the work record (because the first control field is a normal control field) and then changes it if necessary. Notice that any forced character replacements are made *before* the records are sorted.

Force-All

Suppose your control field can contain characters other than A, B, or C. If you don't want a hex FF or 00 in the work and output records, use a force-all line.



Columns 20-22 (Overflow Field Length Entry)

These columns are used only by a summary tag along sort to eliminate the possibility of an overflow condition.

You can eliminate the possibility of an overflow condition by increasing the length of the summary data field to allow for any anticipated overflow. This can be done by coding an overflow field length entry in columns 20-22 of the field description specifications. The overflow field length entry should reflect the sum of the summary data field length and the anticipated overflow length. The entry must be right-justified to column 22, and cannot exceed the maximum field length.

Columns 20-22 are ignored for an FSZ and FSV since these fields can be only one byte long.

For example, if you want to summarize an unpacked field in positions 6-10 of your input record and you know that the output will exceed the 5-position summary field by 2 positions, specify a 7 in column 22:

0	7	F	S	U			6			1	0					7			
---	---	---	---	---	--	--	---	--	--	---	---	--	--	--	--	---	--	--	--

If packed fields are summarized, columns 20-22 should specify the number of bytes of packed data.

For example, if you want to summarize a packed field in positions 1-3 of your input record and you know that the output will exceed the 3-position packed summary field (5 numbers plus sign) by 1 position, specify a 4 in column 22 (7 numbers plus sign):

0	7	F	S	P			1			3						4			
---	---	---	---	---	--	--	---	--	--	---	--	--	--	--	--	---	--	--	--

Note: When some include sets do not contain summary specifications, care should be taken to align the data for summarization.

Columns 23-39

These columns are not used. Leave them blank.

Columns 40-72 (Comments)

Columns 40-72 are for your comments. If you instructed disk sort to print sequence specifications (column 27 of the header line is either blank or contains a zero), comments are printed along with your sequence specifications. The comments have no effect on the program's operation.

Columns 40-45 are enclosed in dotted lines on the coding sheet, because many programmers like to use these columns to write the names of fields they described in the field lines.

Each job in this section has:

1. An introduction explaining its purpose
2. Filled-out sequence specification sheets
3. Discussion of the sequence specifications

The first six jobs use one or more of the following files and input records. (*Sample Job 7* uses other files and input records, which are explained in the introduction to that job. *Sample Job 8* uses the inventory file and input records, and is included to illustrate multivolume multiple device input. *Sample Job 9* uses the inventory file and input records, and is included to illustrate record selection on hexadecimal input.)

Files	Records in Files	Contents of Records
Inventory file	Inventory records	Quantities of items in stock Number of transactions for each stock item
Transaction file	Issue records	Shipments of items to customers
	Receipt records	Purchases of more items from suppliers (vendors)
	Adjustment records	Corrections to inventory quantities

The format of the four types of records are shown in the following charts:

Record Formats for the Inventory File

Inventory Record Positions	Field Names	Contents ¹
1-6	ITEM	<i>Item information</i> Number of stock item to which the record applies
7-10	CLASS	Class of the item (men's clothing, jewelry, etc)
11-24	DESC	Description of the item
		<i>Stock status and reorder information</i>
25-27	LOC	Location of the warehouse where item is stored
28-29	UNIT	Unit of measure for ordering the item (by the pound, by the dozen, etc)
30-34	COST	Cost of the item per unit
35-39	PRICE	Price (per unit) at which item is sold
40-45	REORD	Reorder point (lowest stock quantity allowed)
46-50	QTY	Number of units to order when reordering
51-56	BAL	Number of units in stock
58-62	ORDER	Number of units ordered, but not yet received
69-74	AVAIL	Number of units available (BAL + ORDER)
75-80	PORDER	Purchase order number
81-86	PDATE	Date of purchase order
87-91	VEND1	First vendor
92-96	VEND2	Second vendor
		<i>Transaction activity</i>
97-99	TRANS	Number of transactions for this period
100-105	TDATE	Data of last transaction
106-111	ISSUES	Number of units issued during this period

¹ Any unsigned numeric field (packed or unpacked) can be treated as character information.

Record Formats for the Transaction File

Record Positions	Field Names	Contents ¹
		<i>Issue Record</i>
1	ID	I (identifies this record as an Issue record)
2-7	ITEM	Number of the stock item to which record applies
8-12	QTY	Number of units ordered (5 dozen, 32 pounds, etc)
13-19	CUST	Customer number
20-27	INV	Invoice number
28-33	DATE	Date of the order
34	DIS	Discount information
35-39	PRICE	Price of the item per unit
40-42	CON1	The constant 001
43-50	COST	Total cost of items ordered
		<i>Receipt Record</i>
1	ID	R (identifies this record as a receipt record)
2-5	CLASS	Class of the item ordered (men's clothing, jewelry, etc)
6-11	ITEM	Number of the item ordered
12-16	PRICE	Price of the item per unit (by the dozen, by the pound, etc)
17-22	PORDER	Purchase order number
23-27	QTY	Number of units ordered
		<i>Adjustment Record</i>
1	ID	A (identifies this record as an adjustment record)
2-7	ITEM	Number of the stock item to which record applies)
8	CODE	Adjustment code (identifies reason for the adjustment)
9-13	QTY	Number of units (dozen, pounds, etc) being added or subtracted
14-17	OK	Authorization for the adjustment
18-23	DATE	Date of the adjustment

¹ Any unsigned numeric field (packed or unpacked) can be treated as character information.

SAMPLE JOB 1: INCLUDE ON FIELD RELATIONSHIP

To illustrate an INCLUDE on field relationship, this sample job produces a file of inventory information for reordering stock items.

The program selects inventory records for stock items for which the quantity on hand (AVAIL field) is less than or equal to the reorder point (REORD field).

Output records are to contain information needed to reorder the items. The information is in these fields: CLASS, ITEM, DESC, UNIT, COST, QTY, AVAIL, PDATE, VEND1, VEND2, and PORDER.

Output records are to be in ascending order by class (CLASS field) and by item numbers (ITEM field) within each class.

Line		Job	Card Match	Sequence (A-D-S)	Stagger Select	Output Record Length	Job Description
01	01	SORTR		10A		64	

Line		Factor 1	Factor 2 (Field, Constant, or Keyword)	Comments
01	01	ID 69	74LEF 40 45	AVAIL FLD LESS OR EQ TO REORD FLD

Line	Field	Reserve	Comments
02	NC 7 10	Total - 10 characters	CLASS CLASS OF ITEM - JEWELRY, ETC
03	NC 1 6		
04	DC 11 24	Total - 54 characters (including VEND1, VEND2, and PORDER on page 02)	ITEM STOCK ITEM NUMBER
05	DC 20 29		DESC DESCRIPTION OF ITEM
06	DC 30 34		UNIT UNIT OF MEASURE
07	DC 46 50		COST COST PER UNIT
08	DC 69 74		QTY REORDER QUANTITY
09	DC 81 86		AVAIL UNITS AVAILABLE - BAL+ORDER
			PDATE PURCHASE ORDER DATE

Since there is not enough space on the first coding sheet for the last three fields, they are coded on a second sheet. To show the proper sequence for the last three fields, put a 02 in columns 1-2 of the second coding sheet:

Line		Job	Card Match		Field		Reserved		Field	
Line Number	Type	Job	From	To	Record Character	Substitute Character	Character	Field Length	Field Name	Field Name
07	DC		87	91					VEND1	FIRST VENDOR
08	DC		92	96					VEND2	SECOND VENDOR
09	DC		75	80					ORDER	PURCHASE ORDER NUMBER

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tag along sort.

Columns 13-17

The control fields used to sort records in this job are CLASS and ITEM. The total of their lengths is 10, which is the entry for columns 13-17.

Column 18

The A stands for ascending order. Output records are to be sorted into ascending order by CLASS and ITEM fields.

Columns 29-32

The output records contain eleven fields, including two control fields (CLASS and ITEM) and nine data fields. Control fields are not being dropped (column 28 blank). Therefore, the length of the output record is the total length (64) of all eleven fields described in field specifications.

Record Type Specifications

Column 6

The I identifies the line as an include line. The input records described by the line are used by the program. All others are omitted.

Column 8

Both factor 1 and factor 2 fields are numeric positive fields. The D in column 8 causes zeros and blanks to be considered equal in case leading zeros were not recorded.

Columns 9-16, 17-18, and 20-27

Inventory records for which the number of units available (AVAIL) is equal to or less than the reorder point (RE-ORD) are selected:

- Columns 9-16 identify the location of the AVAIL field (record positions 69-74).
- Columns 20-27 identify the location of the REORD field (record positions 40-45).
- Columns 17-18 define the relation between the two fields.

Column 19

An F in column 19 indicates that factor 1 is compared to the field located by the entries in columns 20-27.

Field Specifications

Column 6

The F identifies the lines as field specification lines.

Column 7

The first two fields described (CLASS and ITEM) are normal control fields. The rest are data fields. All are included in the output record (column 28 of header line is blank).

The records are sorted into ascending order by the CLASS field. Within each class, the item numbers (ITEM field) are in ascending order.

Column 8

Zone and digit portions of all fields are used as they appear in the input records.

Columns 9-16

Columns 9-16 identify the locations of fields in the input records.

CLASS Sequence	ITEM Sequence	Sequence of Records in Sorted Output File
Lowest class number	Lowest item number Next higher item number Third higher item number . . . Highest item number	Lowest class number -- Lowest item number Lowest class number -- Next higher item number Lowest class number -- Third higher item number . . . Lowest class number -- Highest item number
Each higher class number	Lowest item number Next higher item number Third higher item number . . . Highest item number	Class number -- Lowest item number Class number -- Next higher item number Class number -- Third higher item number . . . Class number -- Highest item number
Highest class number	Lowest item number Next higher item number Third higher item number . . . Highest item number	Highest class number -- Lowest item number Highest class number -- Next higher item number Highest class number -- Third higher item number . . . Highest class number -- Highest item number

SAMPLE JOB 2: ADDRROUT SORT

To illustrate an addrout sort, this sample job produces a record address file for an entire inventory file.

Assume that the inventory file is an indexed file. The ITEM field is the record key used in the index (the record keys are in ascending order).

This job creates a record address file containing relative record numbers of all records in the file. It sorts them into ascending order by the CLASS field. Within each class, addresses are also sorted into ascending order by ITEM field.

As a result of this job, inventory records can be processed consecutively in either of two ways:

1. By using the file index, which is in ascending order by ITEM field.
2. By using the record address file created by the disk sort program, which is in ascending order by CLASS and also by ITEM within each class.

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTA identifies the job as an addrout sort.

Columns 13-17

The control fields used to sort addresses in this job are CLASS and ITEM. The total of their length is 10, which is the entry for columns 13-17.

Column 18

The A stands for ascending order. The records are sorted into ascending order by item class (CLASS field), and by item number (ITEM field) within each class.

Line		Job		Card		Largest size		Card		Largest size		Card		Largest size		Card		Largest size	
Number		Type		Format		Format		Format		Format		Format		Format		Format		Format	
01	H	SORTA		10A															
Line		Factor 1		Factor 2 (Field)		Factor 3 (Field)		Factor 4 (Field)		Factor 5 (Field)		Factor 6 (Field)		Factor 7 (Field)		Factor 8 (Field)		Factor 9 (Field)	
Number		Type		Format		Format		Format		Format		Format		Format		Format		Format	
02																			
03																			
04																			
05																			
06																			
Line		Field		Field		Field		Field		Field		Field		Field		Field		Field	
Number		Type		Format		Format		Format		Format		Format		Format		Format		Format	
07		MC		7		10													
08		ND		1		6													
09																			
10																			
11																			
12																			
13																			
14																			
										Total - 10 characters		CLASS		CLASS OF ITEM--JEWELRY, ETC					
												ITEM		STOCK ITEM NUMBER					

Record Type Specifications

Because all input records are being used and all have the same field specifications, no record type specifications are needed (include-all is implied).

Field Specifications

Column 6

The F identifies the lines as field specification lines.

Column 7

Both fields (CLASS and ITEM) are normal control fields. The record addresses are sorted into ascending order by CLASS field and by ITEM field within each class:

Column 8

The C means that both zone and digit portions of the characters in the CLASS field are used in the sorting process.

The D means that only the digit portions of the characters in the ITEM field are used in the sorting process. This causes blanks to be treated as zeros.

Columns 9-16

The CLASS field is located in positions 7-10 of the input records.

The ITEM field is located in positions 1-6 of the input records.

CLASS Sequence	ITEM Sequence	Sequence of Records in Sorted Output File
Lowest class number	Lowest item number Next higher item number Third higher item number . . Highest item number	Lowest class number -- Lowest item number Lowest class number -- Next higher item number Lowest class number -- Third higher item number . . Lowest class number -- Highest item number
Each higher class number	Lowest item number Next higher item number Third higher item number . . Highest item number	Class number -- Lowest item number Class number -- Next higher item number Class number -- Third higher item number . . Class number -- Highest item number
Highest class number	Lowest item number Next higher item number Third higher item number . . Highest item number	Highest class number -- Lowest item number Highest class number -- Next higher item number Highest class number -- Third higher item number . . Highest class number -- Highest item number

SAMPLE JOB 3: AND LOGIC—OPPOSITE CONTROL FIELDS

The program selects inventory records for stock items in classes 0126-0130.

Output records are to contain only inventory information that indicates the activity (transactions) involving the stock items. The information is in these fields: ITEM, CLASS, TRANS, and ISSUES.

Output records are to be sorted in ascending order by class (CLASS field), and in descending order by number of transactions (TRANS field) within each class.

IBM Information Systems Corporation
SEQUENCE SPECIFICATIONS
 Header Page 011

Line	Code	Function	Field	Length	Position	Comments
1	NC	7	10	7	10	} Total — 7 characters
2	OD	97	99	2	99	
3	DC	1	6	1	6	} Total — 19 characters
4	DC	97	99	2	99	
5	DC	7	10	7	10	
6	DC	106	111	6	111	

Watch sequence 095

CLASS CLASS OF ITEM--JEWELRY, ETC
 TRANS NUMBER OF TRANSACTIONS
 ITEM STOCK ITEM NUMBER
 TRANS NUMBER OF TRANSACTIONS
 CLASS CLASS OF ITEM--JEWELRY, ETC
 ISSUES NUMBER OF UNITS ISSUED

Record Type
 INCLUDE CLASS FLDS 0126-0130

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tag along sort.

Columns 13-17

The control fields used to sort records in this job are CLASS and TRANS. The total of their lengths is 7, which is the entry for columns 13-17.

Column 18

The A stands for ascending order. One control field (CLASS) is defined as a normal control field in field specifications. The other control field (TRANS) is an opposite control field. The records, therefore, are sorted into ascending order by item class (CLASS field), and into descending order by number or transactions (TRANS field) within each class.

Column 28

The X indicates that control fields are being dropped from the output records because the TRANS field is an opposite control field. The program must convert the information from the TRANS field into a special form so that the records can be sorted properly. The program does not return the information to its original form after sorting. The information, therefore, would be meaningless.

The information from these two fields, however, is necessary in the output records. Therefore, the fields are also defined as data fields (see the explanation of column 7 of the field specifications for this example).

Columns 29-32

Because control fields are dropped from the output records, the length of the output records is the total length of the fields defined as data fields in field specifications. The length is 19, which is the entry in columns 29-32.

Record Type Specifications

Column 6

The I identifies the record type lines as include lines. The input records described by these lines are used by the program. All others are omitted.

Column 7

The A in column 7 means that line 02 is a continuation of the definition of the record described in line 01. A means AND. An input record must meet all conditions listed in lines 01 and 02 before it is included.

Column 8

The D in column 8 of lines 01 and 02 means that the digit portions of the characters in positions 7-10 of the records (factor 1) is compared with the digit portion of the constants 0126 and 0130 (factor 2). Leading blanks in positions 7-10 would be treated like zeros. Therefore, \emptyset 126 (\emptyset is blank) and 0126 would be considered equal.

Columns 9-16, 17-18, and 20-39

Inventory records that indicate the activity (transactions) of stock classes 0126-0130 are to be selected:

- Columns 9-16 (factor 1) identify the location of the CLASS field (record positions 7-10). The CLASS field contains 4-digit class numbers.
- Columns 20-23 list constants (C in column 19), that in this example define the upper and lower limits of the classes being selected from the inventory file.
- Columns 17-18 define the relationship between the input record field being used to build the file (factor 1) and the constants in factor 2. All inventory classes equal to or greater than (GE) 0126 and equal to or less than (LE) 0130 are to be included in the sorted file.

Columns 19 and Columns 20-39

The C in column 19 indicates that factor 1 is compared to a constant in columns 20-39. The constant in line 01 is 0130; the constant in line 02 is 0126.

Field Specifications

Column 6

The F in column 6 of lines 07-12 indicates that the lines are field specification lines.


Column 7

Column 7 indicates the type of fields being described. The fields described in lines 07 and 08 (CLASS and TRANS, respectively) are control fields. CLASS is a normal control field. TRANS is an opposite control field. Output records are sorted so that CLASS fields are in the order indicated in column 18 of the header specifications (ascending order). Records having the same CLASS field are in descending order according to the TRANS field.

The fields described in lines 09-12 are data fields. They are included in the output record but have no effect on the sequence of the records.

Note that the CLASS and TRANS fields are described once as control fields and once as data fields because TRANS is an opposite control field. Its contents in the control field portion of the output records are changed in the sorting process and would be meaningless for later processing. Therefore the control fields are dropped from the output records (X in column 28 of the header specifications) and CLASS and TRANS are defined as data fields so they would still appear in the output records.

The form of the output record is as follows:



SAMPLE JOB 4: MULTIPLE INCLUDE SETS WITH FORCE

This sample job produces a history file of transaction information by selecting records for all transactions in the transaction file.

The entire file of input records is used in creating output records. The output records will be used to produce a history of the activity involving various stock items.

Output records are sorted by record type: RECEIPT records first, ISSUE records second, and ADJUST (adjustment) records last. Each type of record is sorted into ascending order by item number (ITEM field). The ISSUE records also have a subcontrol field – DATE OF ORDER.

IBM	SEQUENCE SPECIFICATIONS	Header	Page 02	Printed in US-A
Date	Job	Card Method	Job Description	Job Description
Number	Job	Number	Job Description	Job Description
	<i>SORTR</i>	<i>13A</i>	<i>39</i>	
Record Type	Field	Comments		
Item	Field	Comments		
<i>I</i>	<i>C</i>	<i>IEQR</i>	<i>RECEIPT</i>	
FC	6	11	Total -- 7 characters Total -- 20 characters	FORCED CONTROL FIELD NUMBER OF ITEM ORDERED CLASS CLASS OF ITEM--JEWELRY, ETC PRICE PRICE OF ITEM PER UNIT ORDER PURCHASE ORDER NUMBER QTY NUMBER OF UNITS ORDERED
WD	2	5		
DC	12	16		
DC	17	22		
DC	23	27		

IBM

SEQUENCE SPECIFICATIONS

Header

Page 02

Line	Job	Card Match	Record Type	Comments
01	I	C	IEQCI	ISSUE

Line	Number	Type	Location	Record Character	Substitute Character	Overflow Field Length	Record	Field Name	Comments
01	FC		2	7			Total - 13 characters	ITEM	FORCED CONTROL FIELD
02	ND		8	12		Total - 26 characters			
03	DC		13	19			Total - 26 characters	RTY	DATE OF ORDER
04	DC		20	27		Total - 26 characters			
05	DC		34	39			Total - 26 characters	INV	CUSTOMER NUMBER
06	DC		35	39		Total - 26 characters			
07	DC						Total - 26 characters	PRICE	DISCOUNT INFORMATION
08	DC					Total - 26 characters			

IBM

SEQUENCE SPECIFICATIONS

Header

Page 03

Line	Job	Card Match	Record Type	Comments
01	I	C	IEQCA	ADJUST

Line	Number	Type	Location	Record Character	Substitute Character	Overflow Field Length	Record	Field Name	Comments
01	FC		2	7			Total - 7 characters	ITEM	FORCED CONTROL FIELD
02	ND		8	13		Total - 16 characters			
03	DC		9	13			Total - 16 characters	QTY	ADJUSTMENT CODE
04	DC		14	17		Total - 16 characters			
05	DC		18	23			Total - 16 characters	DATE	AUTHORIZATION
06	DC					Total - 16 characters			

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tag along sort.

Columns 13-17

Three types of records are described, one on each sheet. The total length of the control fields for each type of record is 7, 13 and 7 characters, respectively. The largest of these totals (13) is the entry used in columns 13-17.

Column 18

The A stands for ascending order. It indicates the order into which output records (described in the three sets of field lines) will be sorted.

Columns 29-32

The total length of control fields for each type of record described is 7, 13, and 7 respectively. The total length of data fields for each type of record is 20, 26, and 16 respectively. The largest control field total (13) is added to the largest data field total (26) and the result (39) is used in columns 29-32.

Record Type Specifications

Column 6

Column 6 in line 01 of each of the three sheets contains an I indicating that the lines are include lines. The input records described by these lines are used by the program. All others are omitted.

Column 8, Columns 9-16, Columns 17-18, Column 19, and Columns 20-39

Column 8 in line 01 of each of the three sheets contains a C indicating that the *entire character* in position 1 of the input records is to be compared with the constants R, I, and A respectively. Only input records with one of those characters in position 1 will be included in this job.

In the record type lines, columns 9-39 have the following meanings:

- Columns 9-16 identify the input record position (position 1) being compared against R, I, and A.
- Columns 17-18 indicate the comparison result that determines whether an input record is used (the character in position 1 must *equal* R, I, or A).
- C in column 19 of each line indicates that the characters against which the input record character is compared are constants supplied in columns 20-39.
- Columns 20-39 supply the constants (R, I, and A, respectively) used in the comparison.

Field Specifications

Field specifications differ for the three types of records being described. The records, therefore, were described on separate coding sheets.




Column 6

The F in column 6 of lines 07-12 on sheet 01, lines 07-14 on sheet 02, and lines 07-12 on sheet 03 indicate that the lines are field lines.

Column 7

Column 7 indicates the types of fields being described for each of the three types of records. The records described by sheets 01 and 03 (RECEIPT and ADJUST records respectively) have two control fields. The records described by sheet 02 (ISSUE records) have three. The remaining fields described are data fields.

These diagrams show the formats of the output records that will be created from each of the three types of input records. The numbers underneath the field identification blocks indicate the position that information will occupy in the output record.



The first control field in each type of output record is a forced control field. The number 1 will begin every RECEIPT output record, 2 will begin ISSUE records, and 3 will begin ADJUST records. This forces all RECEIPT output records to be first, all ISSUE records second, and all ADJUST records third.

The output records for this job must be the same length. The total number of characters in the control fields described for RECEIPT and ADJUST was less than the total for ISSUE records (the control field length for ISSUE on the header is used to assign the absolute length of the control field). Therefore, the program added binary zeros at the end of the control fields for RECEIPT and ADJUST output records. Blanks were added at the end of the data portion of RECEIPT and ADJUST output records to make the total length of the records equal to the total length of ISSUE output records.

Columns 9-16 and Column 18

Columns 9-16 of the field lines identify the locations of the fields in the input records.

The first control field for each type of record is a forced control field. On each sheet, it is described by a force-all line. For this reason, columns 9-16 are blank in those lines. Column 18 identifies the character to be used in the forced control field: 1 in output records created from RECEIPT records, 2 in ISSUE output records, and 3 in ADJUST output records.

SAMPLE JOB 5: SUMMARY SORT—ELIMINATING DUPLICATES

This sample job produces a list of inventory items requiring adjustments by selecting all adjustment records in the transaction file.

The output record is a control field consisting of the stock item number only.

Since the output record consists of no data field and, in particular, no summary data fields, the result of the SORTRS job will be to eliminate all duplicate adjustment stock item numbers from the file.

The output records, each consisting of a unique stock item number that had an adjustment, are to be in ascending order.

The output provides a list of all stock items that had an adjustment.

Line		Job	Card Mark		Output Record Length		Job Description	
Number	Type	Card Mark	Card Mark	Card Mark	Card Mark	Card Mark	Card Mark	Card Mark
000	00	SORTRS	6A					

Line	Number	Type	Field 1	Field 2	Field 3	Field 4	Field 5	Field 6	Field 7	Field 8	Field 9	Field 10	Field 11	Field 12	Field 13	Field 14	Field 15	Field 16	Field 17	Field 18	Field 19	Field 20	Field 21	Field 22
01	01	I	EQCA																					
02	02	E																						

Line	Number	Type	Field 1	Field 2	Field 3	Field 4	Field 5	Field 6	Field 7	Field 8	Field 9	Field 10	Field 11	Field 12	Field 13	Field 14	Field 15	Field 16	Field 17	Field 18	Field 19	Field 20	Field 21	Field 22
07	07	MC	2	7																				
08	08																							

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTRS identifies the job as a summary sort.

Columns 13-17

Only one control field is used in this job: ITEM. Its length is 6, which is the entry for columns 13-17.

Column 18

The A stands for ascending order.

Columns 29-32

The output record consists of one control field only. The output record length is 6, the control field length.

Record Type Specifications

Column 6

The I identifies the record type line 01 as an include line. Input records described by this line are to be sorted. Since this is the only include set, all other records are omitted.

Column 7

Line 01 is the first (and only) include line of this include set consisting of lines 01 and 07. Therefore, column 7 must be blank.

Column 8

The C in column 8 tells disk sort to use both the zone and digit portions of the factor 1 and factor 2 fields described in line 01.

Columns 13-16, 17-18, 19, and 20

Only adjustment records containing an A in position 1 are to be included in the sort.

Field Specifications

Column 6

The F in column 6 of line 07 indicates that this is a field specification line.

Columns 7 and 8

The NC in columns 7 and 8 of line 07 indicates that it is a normal control field consisting of both the zone and digit portions of the characters in the input records included in the sort.

Columns 9-16

Columns 9-16 of line 07 identify the location of the stock item number (having a length of 6) that is to be used as the control field.

SAMPLE JOB 6: SUMMARY SORT—INSERTION OF CHARACTER DATA—FDV

This sample job produces a summary history file of issues by customer number by selecting all ISSUE records from the transaction file.

Using a summary sort (SORTRS), a file is to be formed consisting of each customer to whom any stock item was issued.

Each output record consists of a unique customer number, the total number of issues it had, and the total dollar value of all the issues.

The output records are to be in ascending order by customer number.

Line		Job	Card Match		Stacker Select		Output Record Length		CCP Disk Sort Only		Job Description	
Number	Type	Card MATCH, SORT SELECT, PAGE	Largest sum (in bytes) of Control Field Length for Any Record Type	S	M	P	U	S	P	U	Record Length	
01	H	SORTRS	7A								23	

Line	Type of G	Factor 1	Rel	Factor 2 (Field, Constant, or Keyword)	Comments
01	I	C	EQ	IEQLI	ISSUE

Line	Type	Location	Record Character	Overflow Field Length	Record	Field Name	Comments
01	H	13	19		Total - 7 characters	CUST	
02	H						BLANK SPACE
03	H						BLANK SPACE
04	H						BLANK SPACE
05	H						BLANK SPACE
06	H						BLANK SPACE
07	H						BLANK SPACE
08	H						BLANK SPACE
09	H						BLANK SPACE
10	H						BLANK SPACE
11	H	40	42		Total - 16 characters	CONI	TOTAL # OF ISSUES
12	H						BLANK SPACE
13	H						BLANK SPACE
14	H						BLANK SPACE
15	H						BLANK SPACE
16	H	43	50			COST	TOTAL \$ OF ISSUES
17	H						OVERFLOW INDICATOR IS X

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTRS identifies the job as a summary sort.

Columns 13-17

The work record as well as the output record consists of a control field made up of the customer number. Its length is 7, which is the entry for columns 13-17.

Column 8

The A indicates that the control field is to be sorted in ascending order.

Columns 29-32

The control field is not to be dropped when writing the output file. The length of the output record is the control field (seven characters) added to the data fields specified by lines 08-14 (16). This totals 23 and is the entry for columns 29-32.

Record Type Specifications

Column 6

The I identifies the record type line 01 as an include line. Input records described by this line are to be sorted. Since this is the only include set, all others are omitted.

Column 7

Line 01 is the first (and only) include line of this include set consisting of lines 01 and 07; therefore, column 7 must be blank.

Columns 8, 16, 17-18, 19, and 20

Only ISSUE records containing the full character I in position 1 of the input records are to be included in the sort.

Field Specifications

Column 6

The F in column 6 of lines 07-14 indicates that the lines are field lines.

Columns 7 and 8

The NC in line 07 specifies that this is the control field to be sorted in the sequence specified in column 18 of the header line (ascending order) using both the zone and digit portions of the characters.

Lines 08, 09, 11, and 12 are data summary overflow fields which are used here as data field unconditional 1-character force fields.

Line 10 indicates that this is a summary data field. When the data in the input record (positions 40-42) is used to build a work record, only the digit portion of the characters is used. This is done since CON1 is an unsigned decimal number. Using a D in column 8 with signed number causes the absolute value of that value to be used.

Line 13 also is specified as a summary data field. However, U in column 8 indicates that this is signed unpacked decimal data. The sign is to be moved to the work record and used in subsequent summary additions.

Line 14 is specified as the summary overflow indicator byte.

Columns 9-16

Entries are placed in the from and to columns for lines 07, 10, and 13 to indicate the control field, digit summary data field (CON1), and the signed unpacked decimal summary field cost, respectively.

For lines 08, 09, 11, and 12, the entries are blank since these are overflow indicator specifications.

Column 17


The X in line 14 defines the field as an overflow field.

Column 18

This column is only used for lines 08, 09, 11, 12 and 14.
 The line numbers and the forced value in the data fields are:


- 08 ␣ (blank)
- 09 #
- 11 ␣ (blank)
- 12 \$
- 14 ␣ (blank)

The initial work record created from the ISSUE records will look like this:



Since this is a summary sort, the CON1 and COST fields will be summarized for records having identical control fields (CUST). These positions 10-12 (CON1) in the work records will be summarized to reflect the total number of issues (TNUMB) for a given customer number. Likewise, positions 15-22 (COST) will be summarized to reflect the total cost of all issues (TCOST).

The output record will look like this:



A blank in position 23 indicates that TNUMB and TCOST are valid sums. X in position 23 indicates that overflow occurred summarizing TNUMB and/or TCOST for this customer number giving invalid sums.

SAMPLE JOB 7: NEGATIVE VALUES POSSIBLE

This sample job illustrates the possibility of ending up with negative values by sorting the summary history file by descending total dollars issued.

The input for this job is the output file created by *Sample Job 6*. Each input record has the following format:

Record Position	Field Name	Contents
1-7	CUST	Customer number
8-9		Constants: $\text{\textcircled{b}}$ and $\#$
10-12	TNUMB	Total number of issues to customer
13-14		Constants: $\text{\textcircled{b}}$ and $\$$
15-22	TCOST	Total cost of all issues to customer
23	OFLOW	Overflow indicator: $\text{\textcircled{b}}$ — TNUMB and TCOST are valid sums $\text{\textcircled{X}}$ — Invalid sum; overflow occurred while summarizing TNUMB and/or TCOST

Only records with valid TNUMB and TCOST will be used.

A record sort of the valid records is desired; that is, the format of the output record is to be identical to the input (the output of the second summary sort sample job).

The records are to be sorted by total dollars issued (TCOST) in descending order. It is possible that TCOST could be negative for some customer numbers. For this reason, a sort that considers the sign must occur.

Header Specifications

Column: 6

The H identifies the line as a header line.

Columns 7-12

The A record sort involving no summary data required a tag along sort. SORTR is the chosen entry.

Columns 13-17

One type of record is specified to be included. The total length of the control fields for this type is 8 (the entry in column 17).

Column 18

The D stands for descending order. The largest positive number comes first, the largest negative number (the negative number having the largest absolute value) comes last.

The absolute values of -19 and -99 are 19 and 99. -99 is more negative than -19 since its absolute value is greater than 19.

Column 28

The X indicates that the control field is to be dropped when forming the output file.

Columns 29-32

When the control field is dropped, only the data portion remains. This data portion is the entire input record; thus, 23 is entered in columns 29-32.

Record Type Specifications

Columns 6 and 7

Two record types are specified: One type is to be omitted (O in column 6); one is to be included. Each record type is specified by one record type line. Since these lines are the first (as well as the last) record type lines in their sets, column 7 must be blank.

Columns 8, 9-16, 17-18, 19, and 20-39

For statement 0101, these columns specify that if position 23 in the input record contains the full character constant X, the record is to be omitted.

Statement 02 includes all records not omitted by statement 01. For each record included, a work record is formed by the field specifications, statements 07 and 08. Note that the control field is an 8-position unpacked field (statement 07).

Columns 7, 8, 9-16, and 18

Statement 07 indicates unpacked data in the control field from positions 15 through 22 of the input record. (This control field is dropped because an X has been coded in column 28 of the header specifications.)

For a more complete discussion of numeric data see Column 8 and Columns 9-16 under Record Type Specifications.

Statement 08 indicates that the entire record is included as data.

Field Specifications

Column 6

The F in column 6 of statements 07 and 08 indicates that these are field lines.

Line		Statement	Field	Location	Length	Character	Control	Comments
IBM SEQUENCE SPECIFICATIONS								
Header								
1	01	01	SO	15	8	D	X	23
Record Type								
1	02	02	OC	15	23	EQCX		INVALID SUMS INCLUDE-ALL POSITIVE TCOST
Field								
1	07	07	NU	15	22			NORMAL UNPACKED
1	08	08	DC	1	23			ENTIRE RECORD-DATA FIELD

SAMPLE JOB 8: MULTIPLE INPUT FILES

This sample job illustrates the use of multiple input files by producing a 5-year history of sales of selected classes of items from selected vendors.

The input is from the current year inventory file (multi-volume file) and from the inventory files for the past four years (single-volume files on four input devices).

The record format for all files is the same as that used for the inventory files.

The records are selected by CLASS number 0126 and VENDOR numbers 00023, 00108, and 01120.

Output records contain information needed to identify which vendor's line of this CLASS has sold the best (number of transactions and number of units) per year for the current year and the past four years. The information is in the following fields: ITEM, CLASS, VEND1, TRANS, TDATE, ISSUES.

Output records are to be sorted in ascending order by year (last two positions of TDATE field), then by VENDOR number (VEND1 field, and finally by ITEM number (ITEM field).

OCL required for the program is given in the following figure.

```

IBM                                     IBM System/3 Basic Assembler Coding Format
PROGRAM
PROGRAMMER
NAME
OPERATION
OPERAND
SUBJECT
//E
// LOAD $DSORT, F1
// FILE NAME-INPUT1, UNIT- 'R1, F2, R2', PACK- 'R1R1R1, F2F2F2, R2R2R2', RETAIN-P, LABEL-INVENTORY
// FILE NAME-INPUT2, UNIT-T1, REEL-DDDDDD, LABEL-INVR04, RETAIN-999, BLKL-1110, RECL-111, RECFM-FB
// FILE NAME-INPUT3, UNIT-T2, REEL-CCCCC, LABEL-INVR03, RETAIN-999, BLKL-1110, RECL-111, RECFM-FB
// FILE NAME-INPUT4, UNIT-T3, REEL-BBBBBB, LABEL-INVR02, RETAIN-999, BLKL-1110, RECL-111, RECFM-FB
// FILE NAME-INPUT5, UNIT-T4, REEL-AAAAAA, LABEL-INVR01, RETAIN-999, BLKL-1110, RECL-111, RECFM-FB
// FILE NAME-WORK, UNIT-F1, PACK-F1F1F1, TRACKS-20, RETAIN-S
// FILE NAME-OUTPUT, UNIT-R1, PACK-R1R1R1, TRACKS-20, RETAIN-T
// RUN
    
```

SEQUENCE SPECIFICATIONS
Header

01

1021 9089 3
11 11 11 11 11
11 11 11 11 11

Line	Type	Code	Location	Record Length	Record Type	Field	Subfield	Value
009	I	SORTR	13A					26

Record Type

Line	Type	Code	Location	Record Length	Record Type	Field	Subfield	Value
01	I	D	7	10EQC0126				
02	I	A	87	91EQC00023				
03	I	D	7	10EQC0126				
04	I	A	87	91EQC00108				
05	I	D	7	10EQC0126				
06	I	A	87	91EQC01120				

INCLUDE CLASS 0126 AND
VENDOR 00023 OR ALSO INCLUDE
CLASS 0126 AND
VENDOR 00108 OR ALSO INCLUDE
CLASS 0126 AND
VENDOR 01120

Field

Line	Type	Code	Location	Record Length	Record Type	Field	Subfield	Value
07	F	NC	104	105				
08	F	NC	87	91				
09	F	NC	1	6				
10	F	DC	7	10				
11	F	DC	97	99				
12	F	DC	106	111				

DATE YEAR POSITIONS OF DATE
VEND1 VENDOR NUMBER
ITEM STOCK ITEM NUMBER
CLASS CLASS OF ITEM
TRANS NUMBER OF TRANSACTIONS
ISSUES NUMBER OF UNITS ISSUED

Header Specifications

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tag along sort.

Columns 13-17

The control fields used to sort records in this job are the last two positions of the TDATE field, and VEND1 and ITEM. The total of their lengths is 13, which is the entry for columns 13-17.

Column 18

The A stands for ascending order. Output records are to be sorted in ascending order by YEAR, VENDOR number, and ITEM number.

Columns 29-32

The output records contain six fields, including the three control fields. Control fields are not being dropped (column 28 blank). The output length is the total length (26) of all six fields described in field specifications.

Record Type Specifications

Column 6

The I identifies the line as an include line. The input records that match the specifications of these lines are included. All other records are omitted.

Column 7

The A identifies the line as an AND statement while O identifies the line as an OR statement. The sequence of the lines is important in determining the include logic. The logic shown includes three types of input records, each type being included only if the two AND include requirements are met.

Column 8

The D means that the digit portions of the characters in positions 7-10 and 87-91 of the records (factor 1) is compared with the digit portion of the constants (factor 2). Leading blanks in positions 7-10 would be treated like zeros.

Columns 9-16, 17-18, and 20-39

Inventory records that indicate the activity of stock CLASS equal to the constant 0126 and vendor (VEND1 field) equal to the constants 00023, 00118, and 01120 are to be selected.

Column 19 and 20-39

The C in column 19 indicates that factor 1 is compared to a constant in columns 20-39.

Field Specifications

Column 6

The F in column 6 of lines 7-12 indicates that the lines are field specification lines.

Column 7

Indicates the type of fields being described. The fields described in lines 07, 08, and 09 (TDATE, VEND1, and ITEM) are control fields. N specifies they are normal control fields so that output records will be sorted in the order indicated in column 18 of the header specifications.

The fields described in lines 10-12 are data fields. The form of the output record is as follows:

	TDATE	VEND1	ITEM	CLASS	TRANS	ISSUES
Positions	1 ↔ 2	3 ↔ 7	8 ↔ 13	14 ↔ 17	18 ↔ 20	21 ↔ 26

Column 8

The C defines both the zone and digit portion of the input record used.

Statement 09 defines only the digit portion of the input record used.

Columns 9-16

Identifies the location of the fields in the input record.

SAMPLE JOB 9: RECORD SELECTION ON BINARY VALUES

Data selection on a binary value can be accomplished using the System/3 96-column card, 64-character set. Two types of include sequences can be used to accomplish binary selection.

Type 1 Include Sequence

For binary values whose hexadecimal equivalents can be represented by one or two characters in the 64-character set (see Appendix C, *Collating Sequence*), use either an *include character* (I and C in columns 6 and 8 of the record type specifications) or an *include zone* (I and Z in columns 6 and 8 of the record type specifications) ANDed with an *include digit* (IAD in columns 6 through 8 of the record type specifications).

Type 2 Include Sequence

For binary values whose hexadecimal zone portions do not appear in the 64-character set, use two include specifications for the zone portion ANDed with an *include digit* for the digit portion.

Example: Select records that contain a value of 44 in position 2. The binary representation of this number is 0010 1100, which is a hex 2C. For the digit portion, select a character from the table in Appendix C with a digit portion of C (for example, \$, *, %, and @). Use the character in an include line, comparing its digit portion to the binary value of 1100. Since there is no character in the 64-character set whose zone portion is equivalent to 2, use two include lines to describe constants with a zone portion greater than 1 and less than or equal to 2. Describe these constants as packed, which reverses the zone and digit portions from hex F1 and F2 to hex 1F and 2F.

Record Type

Line	Field	Operator	Value	Record	Comments
01	I D		2EQC@		DIGIT = C AND ≤ HEX 2F AND > HEX 1F
02	IAP		2LEC2		
03	IAP		2GTC1		
04					

	Hexadecimal Value		Binary Value		Decimal Value
Field to be selected	2C	=	0010 1100	=	44
Character @	7C	=	0111 1100	=	124
Packed character 2	2F	=	0010 1111	=	47
Packed character 1	1F	=	0001 1111	=	31

The record type specifications shown will cause selection of all records containing (in column 2) a value between 31 (hex 1F) and 47 (hex 2F) whose digit portion is equal to C.

Note: The 32 hexadecimal values whose zone portions are either A or B cannot be accessed with these types of entries. Therefore, there can be no record selection based on these values.

Sample Sort Job 9 Description

The purpose of this sample job is to select records from an inventory file whose reorder point is 7800. The reorder point is a 2-byte quantity in columns 44 and 45 of the inventory records. The quantity is stored in binary (7800 = 0001 1110 0111 1000), and the hexadecimal representation of this number is 1E78. (Use Appendix C to help with binary and hexadecimal conversions.)

SEQUENCE SPECIFICATIONS

Header

Line	Type	Sequence	Start	End	Length	Control	Record	Job Description
01	H	SORTR	6A					

Record Type

Line	Type	Factor	From	To	Character	Constant	Comments
01	I	D			44EQC+		DIGIT = 'E' **
02	I	A			44LEC1		AND <= 'IF' ** TYPE II
03	I	A			44GTC0		AND > '0F' **
04	I	***					*** AND *** **
05	I	A			45EQC@		ZONE = '7' **
06	I	A			45EQC8		DIGIT = '8' **

Field

Line	Type	Location	Character	Constant	Comments
07	I	ND	7	111	SEQUENCE BY ITEM PICK UP REMAINDER OF THE RECORD

	Hexadecimal Value	Binary Value	Decimal Value
Field to be selected, column 44	1E	= 0001 1110	= 30
Character +	4E	= 0100 1110	= 78
Packed character 1	1F	= 0001 1111	= 31
Packed character 0	0F	= 0000 1111	= 15
Field to be selected, column 45	78	= 0111 1000	
Character @	7C	= 0111 1100	
Character 8	F8	= 1111 1000	

IBM

International Business Machines Corporation

G-21-0889
Printed in U.S.A.

SEQUENCE SPECIFICATIONS
Header

Page 2
Program Identification 75 76 77 78 79 80

Line Number	Job	Job Description	Card Match										Output Record Length	Respect	Job Description	
			Stacker Select													
			S	M	P	L	P	P	P	P	P	P	P	P	P	P
01	SDRTR															

Record Type

Line Number	Job	Location	Field	Factor 2 Field or Constant		Comments
				From	To	
01	IAD		44EQCF			DIGIT = 'E' **
02	IAP		44LECI			AND < = 'IF' ** TYPE II
03	IAP		44GTCB			AND > = 'OF' **
04	*** ** *					*** AND ***
05	IAZ		45EQCB			ZONE = '7' ** TYPE I
06	IAD		45EQCB			DIGIT = '8' **

Field

Line Number	Job	Location	Field	Comments
01	WD	1	6	SEQUENCE BY ITEM PICK UP REMAINDER OF THE RECORD
02	DC	7	111	

These specifications will cause selection of all records containing (in column 44) a value between 15 (hex 0F) and 31 (hex 1F) whose digit portion is equal to E. These specifications will further limit selection of records to those which have a binary value in column 45 whose zone portion equals 7 and whose digit portion equals 8. Therefore, all records selected will have hex 1E78 in columns 44 and 45.

Header Specification

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tagalong sort.

Columns 13-17

The total length of the control field is 6 bytes.

Column 18

The A stands for ascending order. The file will be sorted according to the standard collating sequence shown in Appendix C.

Columns 29-32

The control field is not dropped when data is written to the output file (column 28 is blank). Therefore, the length of the output record equals the length of the control field (6 bytes) plus the length of the data field (105 characters) specified by lines 07 and 08. This output record length totals 111, which is the entry for columns 29 through 32.

Record Type Specifications

Line 01 selects records whose digit portion in column 44 is E. Lines 02 and 03 select records whose packed value in column 44 is less than or equal to hex 1F, and at the same time is greater than a packed value of hex 0F.

The only hexadecimal values that satisfy both specifications for lines 02 and 03 are hex 11 through hex 1F. Line 04 is a comment line. Lines 05 and 06 select records whose zone portion in column 45 is 7 and whose digit portion is 8.

Field Specifications

Column 6

The F identifies the lines as field specification lines.

Columns 7, 8, and 9-16

Line 07 indicates that the control field is in positions 1 through 6 of the input records. Only the digit portions of the characters will be used. This control field will be included in the output file.

Line 08 indicates that the other 105 input record characters are included as data in the output file.

SAMPLE JOB 1: SORTING ON A CONTROL FIELD THAT CONTAINS EITHER A POSITIVE INTEGER OR A POSITIVE REAL NUMBER

Line		Job		Card Magn.		Output		Job Description	
Number	Type	Card	Job	Stacker	Printer	Receiver	Length		
000	H	4A	SORTR				64		

Line		Factor 1		Factor 2		Comments	
Number	Type	From	To	From	To		
01	F						
02	F						
03	F						
04	F						
05	F						
06	F						

Line		Field		Comments	
Number	Type	From	To		
07	F	1	4		
08	F	5	64		
09	F				
10	F				
11	F				
12	F				
13	F				
14	F				

Header Specification

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tagalong sort.

Columns 13-17

The control field that contains either a positive integer or a positive real number has a length of 4 bytes (the entry in column 17).

Column 18

The A specifies ascending order. The smallest positive integer or positive real number comes first; the largest positive integer or positive real number comes last in the sort sequence.

Columns 29-32

The control field is not dropped when data is written to the output file (column 28 is blank). Therefore, the length of the output record equals the length of the control field (4 bytes) added to the length of the data field (60 characters) specified by lines 07 and 08. This output record length totals 64, which is the entry for columns 29 through 32.

Record Type Specifications

Because all input records are being used and all have the same field specifications, no record type specifications are needed (include-all is implied).

Field Specifications

Column 6

The F identifies the lines as field specification lines.

Columns 7, 8, and 9-16

Line 07 indicates that the control field is in positions 1 through 4 of the input record. This control field is included in the sorted output file (column 28 of the header line is blank).

SAMPLE JOB 2: SORTING ON A CONTROL FIELD THAT CONTAINS ONLY A POSITIVE OR NEGATIVE INTEGER

Line		Job	Card Match		Output		Job Description	
Number	Code	Code	Marker	Control	Length	Position	Position	Position
00	00	SORTR						
Record Type								
Line	Number	Code	Location	Range	Comments			
01	01	IC			ALL POSITIVE INTEGERS			
Field								
Line	Number	Code	Location	Range	Comments			
02	02	FC			FORCE 2 FOR POSITIVE INTEGERS CONTROL FIELD DATA FIELD			
08	08	NC	1	4				
09	09	DC	1	64				

Record Type Specifications

Column 6

Two record types are specified (positive integers and negative integers). Both record types are to be included (I in column 6 of line 01, pages 01 and 02).

Columns 8, 9-16, 17-18, 19, and 20-39

On line 01 (page 01), these columns specify that if position 1 (the leftmost byte of the 4-byte control field from the input record) is less than or equal to the character constant " (see note), the work record will have a character 2 forced into the first byte of its 5-byte control field (line 07 on page 01). This include line selects only positive integers. Line 01 (page 02), which is an include-all line, indicates that all records not described by the first include line (page 01) are also included in the sort. The first include line (line 01 on page 01) selects all positive integers, therefore, this include-all line (page 2) selects all negative integers, which will have a character 1 forced into the first byte of their work record control field (line 07 on page 02). All negative integers are forced in front of all positive integers in the sort sequence.

Note. The special character " has a bit configuration of 01111111, which is equivalent to the high order byte of the maximum positive integer.

Field Specifications

Column 6

The F identifies the lines on pages 01 and 02 as field specification lines.

Columns 7, 8, 9-16, and 18

On line 07 (pages 01 and 02), these columns specify that for positive integers (line 07 on page 01), a character 2 will be forced into the first byte of the work record control field; for negative integers (line 07 on page 02), a character 1 will be forced into the first byte of the work record control field.

On line 08 and 02, columns 7, 8, and 9 through 16 specify that the control field is in positions 1 through 4 of the input records. (The control field is not included in the output records because header specification column 28, page 01, contains an X.)

SAMPLE JOB 3: SORTING ON A CONTROL FIELD THAT CONTAINS ONLY A POSITIVE OR NEGATIVE REAL NUMBER

Line		Job	Card Match	Card Match		Card Match		Card Match		Card Match		Card Match		Card Match		Card Match		Card Match		Card Match	
Number	Job	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match	Card Match
0 0 0 0	SORTR	5A																			
Record Type																					
Line	Job	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9	Factor 10	Factor 11	Factor 12	Factor 13	Factor 14	Factor 15	Factor 16	Factor 17	Factor 18	Factor 19	Factor 20
0 1	C																				
0 2																					
0 3																					
0 4																					
0 5																					
0 6																					
Field																					
Line	Job	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field	Field
0 7	FC																				
0 8	NC	1	4																		
0 9	DC	1	64																		
1 0																					
1 1																					
1 2																					
1 3																					

Line		Job		Card Match		Sort		Output		CCP		Sort		Job Description	
Number	Type	Code	Code	Code	Code	Code	Code	Code	Code	Code	Code	Code	Code	Code	Code
0	H														

Line		Factor 1		Factor 2 (Field, Constant, or Keyword)		Comments	
Number	Type	Code	Code	Code	Code	Code	Code
0-16	I						
0-1							CATCH-ALL FOR NEG. REAL NUMBERS

Line		Field		Comments	
Number	Type	Code	Code	Code	Code
0-16	I				
0-1		FC	4	1	FORCE 1 FOR NEG. REAL NUMBERS ** OPPOSITE SEQUENCE CONTROL FLD DATA FIELD

Header Specification

Column 6

The H identifies the line as a header line.

Columns 7-12

SORTR identifies the job as a tagalong sort.

Columns 13-17

The total length of the control field (for either a positive or negative real number) is 5 bytes (the entry in column 17).

Column 18

The A stands for ascending order. The smallest negative real number comes first, the largest positive real number comes last in the sort sequence (for example -2.0, -1.0, 0.0, 1.0, 2.0).

Column 28

The X indicates that the control field is dropped when data is written to the output file.

Columns 29-32

When the control field is dropped, only the data portion remains. The data portion in this sample job is the entire input record. Therefore, 64 is entered in columns 29 through 32.

Record Type Specifications

Column 6

Two record types are specified (positive real numbers and negative real numbers). Both types are to be included (1 in column 6 of line 01, pages 01 and 02).

Columns 8, 9-16, 17-18, 19, and 20-39

On line 01 (page 01), these columns specify that if position 1 (the leftmost byte of the 4-byte control field from the input record) is less than or equal to the character constant " (see note), the work record will have a character 2 forced into the first byte of its 5-byte control field (line 07 on page 01). This include line selects only positive real numbers. Line 01 (page 02), which is an include-all line, indicates that all records not described by the first include line (page 01) are also included in the sort. The first include line (line 01 on page 01) selects all positive real numbers; therefore, this include-all line (page 02) selects all negative real numbers, which will have a character 1 forced into the first byte of their work record control field (line 07 on page 02). All negative real numbers are forced in front of all positive real numbers in the sort sequence.

Note: The special character " has a bit configuration of 01111111, which is equivalent to the high order byte of the maximum positive real number.

Field Specifications

Column 6

The F identifies the lines on pages 01 and 02 as field specification lines.

Columns 7, 8, 9-16, and 18

On line 07 (pages 01 and 02), these columns specify that for positive real numbers (line 07 on page 01), a character 2 will be forced into the first byte of the work record control field, for negative real numbers (line 07 on page 02) a character 1 will be forced into the first byte of the work record control field.

On line 08 (pages 01 and 02), columns 7, 8, and 9 through 16 specify that the control field is in positions 1 through 4 of the input records. Note that real numbers are always carried in true form (never in two's complement form). Therefore, sorting negative real numbers into ascending sequence requires them (when treated as character fields) to be sorted into descending sequence.

On line 09 (pages 01 and 02), columns 7, 8, and 9 through 16 specify that 64 input record characters are included as data in the sorted output file.

Appendix A. Conversational OCL for Model 6

The OCL cycle you use depends on the type of sort you want to do:

Type of Sort	Cycle To Use	What The Cycle Does
Sort job you plan to run only a few times	LOAD	Gives the system the OCL statements it needs to run the sort job.
Sort job you plan to run frequently	BUILD	Stores the OCL statements necessary to run the sort job in a source library. ¹
	CALL	Tells the system you want to use a procedure ¹ to run a sort job.
Sort job you plan to run as one of a group of jobs	BUILD	Stores the OCL statements necessary to run the sort job in a source library. ¹
	BUILDC	Puts the name of each job in the group and the disk location of the OCL statements for each job in a source library. ¹
	CALL	Tells the system you want to use the BUILDC procedure ¹ to run a group of jobs.

¹A set of OCL statements in a source library is called a procedure.

Note: See Appendix H for CCP/Disk Sort oriented file and program considerations.

CODING OCL CYCLES

This appendix contains keyword-response summary charts for the LOAD, BUILD, and CALL cycles. For the keyword-response summary for the BUILDC cycle, see the *IBM System/3 Model 6 Operation Control Language and Disk Utility Programs Reference Manual*, GC21-7516.

- Words or letters in all capital letters (such as INPUT, INCLUDE, R1) are entries the operator must make.
- Words or letters not in all capital letters (such as procedure name, mmdyy) represent information you must supply.
- P/S indicates the program start (PROG START) key.
- ENTER- indicates the ENTER - key.

For more information about any particular statement (keyword and response), see the keyword description section of the *IBM System/3 Model 6 Operation Control Language and Disk Utility Programs Reference Manual*, GC21-7516.

THE MODIFY STATEMENT

The only responses shown on the following charts are those that apply directly to the disk sort program. For other MODIFY options (such as correcting one of the statements in a cycle) see the keyword description section of the *IBM System/3 Model 6 Operation Control Language and Disk Utility Programs Reference Manual*, GC21-7516.

THE LOAD CYCLE

Prompt	Response	Considerations
READY	LOAD (P/S)	—
LOAD NAME	\$DSORT (P/S)	System name for the disk sort program
UNIT	R1, R2, F1, or F2 (P/S)	Disk that the disk sort program is on
DATE	mmddyy or ddmmyy (P/S) or (P/S)	If you want a different system date for your job If you want to use the system date for your job or no date is required
SWITCH	(P/S)	Disk sort does not use external indicators
FILE NAME	INPUT (P/S)	Name that disk sort uses for file you want to sort
UNIT	R1, R2, F1, or F2 (P/S)	Which disk the file you want to sort is on
PACK	Disk name (P/S)	Name of disk which contains the file to be sorted
LABEL	VTOC filename (ENTER-)	Name you use for the file you want to sort (PAYROLL, INVENT, etc.)
FILE NAME	WORK (optional) (P/S)	Name disk sort uses to refer to file you are going to use for work area during this sort. If WORK is not used, disk sort performs automatic work file allocation.
UNIT	R1, R2, F1, or F2 (P/S)	Disk containing the file you are going to use for a work area
PACK	Disk name (P/S)	Name of disk which contains the work file
LABEL	(P/S)	—

Input file

THE LOAD CYCLE (continued)

Work file

Output file

Prompt	Response	Considerations
RECORDS	(P/S)	Do not specify RECORDS for the work file
TRACKS	1-398 (P/S)	Number of tracks file occupies (see Appendix E)
LOCATION	8-405 (P/S) or (P/S)	If you want to control the location of the file on the disk If you want to let disk sort control the location of the file on the disk
RETAIN	S (ENTER-)	Scratch status because the work file is used for only one sort job
FILE NAME	OUTPUT (P/S)	Name disk sort uses for the sorted file
UNIT	R1, R2, F1, or F2 (P/S)	The disk you want to put the sorted file on
PACK	Disk name (P/S)	Name of the disk which will contain the sorted file
LABEL	VTOC filename (P/S)	Name you use for the sorted file (this can be the same as your name for the file you want to sort)
RECORDS (see note)	1-999999 (P/S)	Number of records you want to sort
TRACKS (see note)	1-398 (P/S)	Number of tracks file occupies (see Appendix E)
LOCATION	8-405 (P/S) or (P/S)	If you want to control the location of the file on the disk (if you are overlaying an input file with an output file you must specify LOCATION) If you want to let disk sort control the location of the file on the disk
RETAIN	P or T (P/S)	If the output file is built during the sort job, it can have permanent (P) or temporary (T) status. If the output file existed before you started the sort job, the file cannot have permanent status since you cannot overlay a permanent file.

Note: These are the space keywords. If you are creating a new output file you must respond to one of them.

THE LOAD CYCLE (continued)

Prompt	Response	Considerations
FILE NAME	(P/S)	Disk sort uses only three files
MODIFY	RUN (P/S)	Tells disk sort that the OCL statements are complete.
	<ol style="list-style-type: none"> 1. Sequence specifications 2. // END 	<ol style="list-style-type: none"> 1. Operator is going to enter the sequence specifications from the keyboard (one specification per line). 2. Sequence specifications are complete. 3. Disk sort program sorts your file.
	// SOURCE name, unit	<ol style="list-style-type: none"> 1. name = name of the KSE-created source library entry which contains the set of sequence specifications you want to use. unit = the disk (R1, R2, F1, or F2) whose source library contains the sequence specifications you want to use. 2. When all these sequence specifications have been read, the disk sort program sorts your file.

If you are going to enter sequence specifications from the keyboard

If the sequence specifications are in a KSE-created source library entry

THE BUILD CYCLE

Prompt	Response	Considerations
READY	BUILD (P/S)	—
BUILD NAME	Procedure name (P/S)	The name you want to give the procedure: <ul style="list-style-type: none"> ● Maximum of six alphameric characters ● Must begin with alphabetic character ● Must not be DIR, SYSTEM, or ALL
UNIT	R1, R2, F1, or F2 (P/S)	The disk containing the source library you want to put the procedure in
LOAD NAME	SDSORT (P/S)	System name for the disk sort program
UNIT	R1, R2, F1, or F2 (P/S)	Disk that the disk sort program is on
DATE	ddmmyy or mmddyy (P/S) or (P/S)	If you want a different system date for your job If you want to use the system date for your job or no date is required
SWITCH	(P/S)	Disk sort does not use external indicators
FILE NAME	INPUT (P/S)	Name that disk sort uses for file you want to sort
UNIT	R1, R2, F1, or F2 (P/S)	Disk the file you want to sort is on
PACK	Disk name (P/S)	Name of disk that contains the file to be sorted
LABEL	VTOC filename (ENTER-)	Name you use for file you want to sort (PAYROLL, INVENT, etc.)

Input file



THE BUILD CYCLE (continued)

Work file

Prompt	Response	Considerations
FILE NAME	WORK (optional) (P/S)	Name disk sort uses to refer to file you are going to use for work area during this sort. If WORK is not used, disk sort performs automatic work file allocation.
UNIT	R1, R2, F1, or F2 (P/S)	Which disk the file you are going to use for work area is on
PACK	Disk name (P/S)	Name of disk that work file is on
LABEL	(P/S)	—
RECORDS	(P/S)	—
TRACKS	1-398 (P/S)	Number of tracks file occupies (see Appendix E)
LOCATION	8-405 (P/S) or (P/S)	If you want to control the location of the file on the disk If you want to let disk sort control the location of the file on the disk
RETAIN	S (ENTER-)	Scratch status because the work file is used for only one sort job

THE BUILD CYCLE (continued)

Output file

Prompt	Response	Considerations
FILE NAME	OUTPUT (P/S)	Name that disk sort uses for file area where you want to put the sorted file
UNIT	R1, R2, F1, or F2	Disk you want to put the sorted file on
PACK	Disk name (P/S)	Name of disk which will contain the sorted file
LABEL	VTOC filename (P/S)	Name you use for the sorted file (this can be the same as your name for the file you want to sort)
RECORDS (see note)	1-999999 (P/S)	Number of records you want to sort
TRACKS (see note)	1-398 (P/S)	Number of tracks file occupies (see Appendix E)
LOCATION	8-405 (P/S) or (P/S)	If you want to control the location of the file on the disk If you want to let disk sort control the location of the file on the disk
RETAIN	P or T (ENTER-)	If the output file is built during the sort job, it can have permanent (P) or temporary (T) status. If the output file existed before you started the sorted job, the file cannot have permanent status since you cannot overlay a permanent file

Note: These are the space keywords. If you are creating a new output file you must respond to one of them.

THE BUILD CYCLE (continued)

You are going to include sequence specifications in the procedure

You are going to put a pointer to the sequence specifications in the procedure

You are not putting sequence specifications in the procedure

Prompt	Response	Considerations
FILE NAME	(P/S)	Disk sort uses only three files
MODIFY	<ol style="list-style-type: none"> 1. INCLUDE 2. Sequence specifications 3. // END 	<ol style="list-style-type: none"> 1. Tells the system you are going to include sequence specifications in the procedure 2. Operator types the specifications (one specification per line) 3. Tells the system the sequence specifications are complete
MODIFY	RUN (P/S)	System gives operator opportunity to correct any mistakes in sequence specifications. Operator types RUN when sure all the specifications are correct. System stores the procedure in the source library you specified in the second OCL statement.
MODIFY	<ol style="list-style-type: none"> 1. INCLUDE 2. // SOURCE name, unit 	<ol style="list-style-type: none"> 1. Tells the system you are going to include in the procedure a pointer to the sequence specifications 2. name = name of the KSE-created source library entry which contains the set of sequence specifications you want to include in the procedure unit = the disk (R1, R2, F1, or F2) whose source library contains the set of sequence specifications you want to include in the procedure
MODIFY	RUN (P/S)	You are not going to include sequence specifications in the procedure. The system stores the procedure in the source library you specified in the second OCL statement.

THE CALL CYCLE

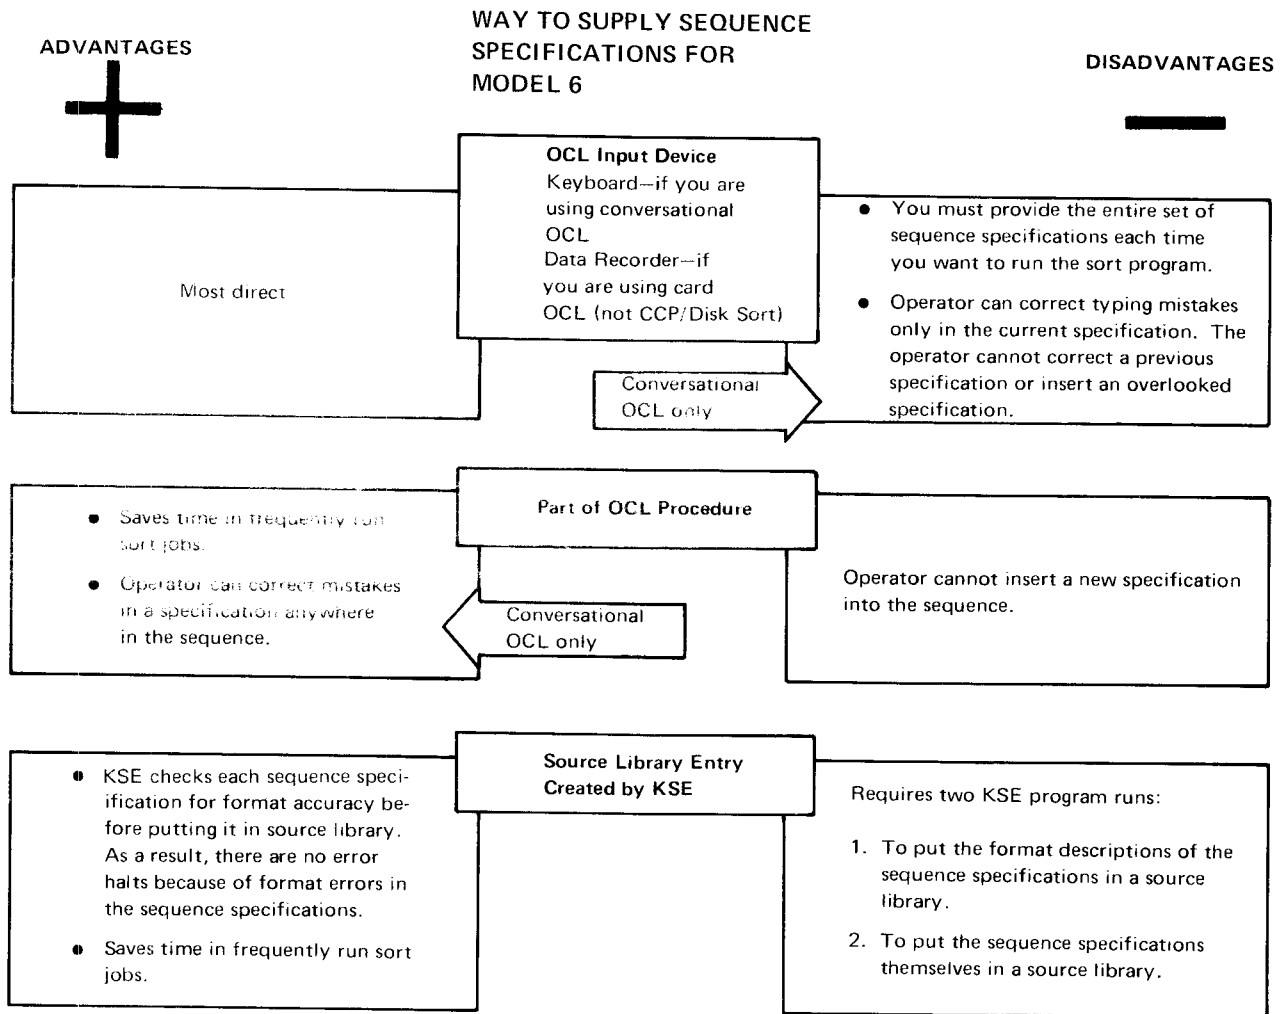
Prompt	Response	Considerations
READY	CALL (P/S)	
CALL NAME	Procedure name (P/S)	Your response to BUILD NAME during the BUILD sequence
UNIT	R1, R2, F1, or F2 (P/S)	Disk where the procedure is located; system displays the procedure
MODIFY	RUN (P/S)	Tells the system the OCL statements are correct
INCLUDED STATEMENTS	No response	System prints sequence specifications
MODIFY	RUN (P/S)	Tells the system the sequence specifications are correct; system runs the sort job
MESSAGE SD100 ASKING FOR SEQUENCE SPECIFICATIONS	1. Sequence specifications 2. // END	1. Operator types sequence specifications on the keyboard (one specification per line) 2. Tells the system the sequence specifications have been entered
MESSAGE SD100 ASKING FOR SEQUENCE SPECIFICATIONS	// SOURCE name, unit	name = name of the KSE-created source library entry which contains the sequence specifications you want to use. unit = location of disk whose source library entry you want to use (R1, R2, F1, or F2) When the system has read all the sequence specifications, it runs the sort job

If sequence specifications are included in the procedure

If sequence specifications are not in the procedure and you want to enter them through the keyboard

If sequence specifications are not in the procedure but are in a KSE-created source library entry

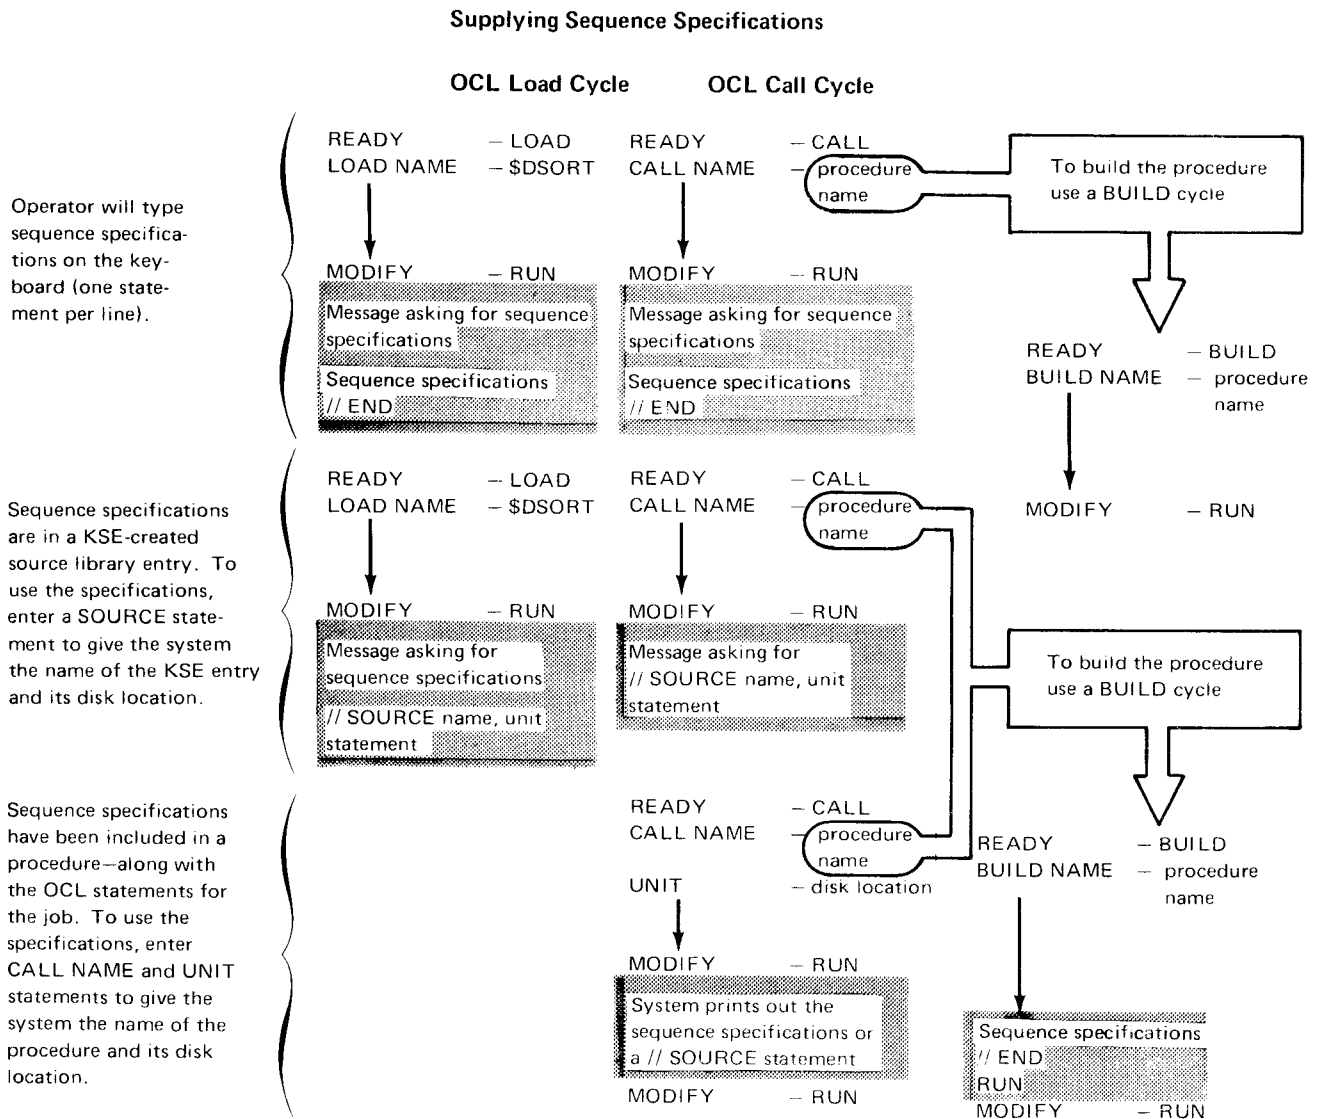
Appendix B. Supplying Sequence Specifications



INTERRELATIONSHIP OF CONVERSATIONAL OCL AND SEQUENCE SPECIFICATIONS (MODEL 6)

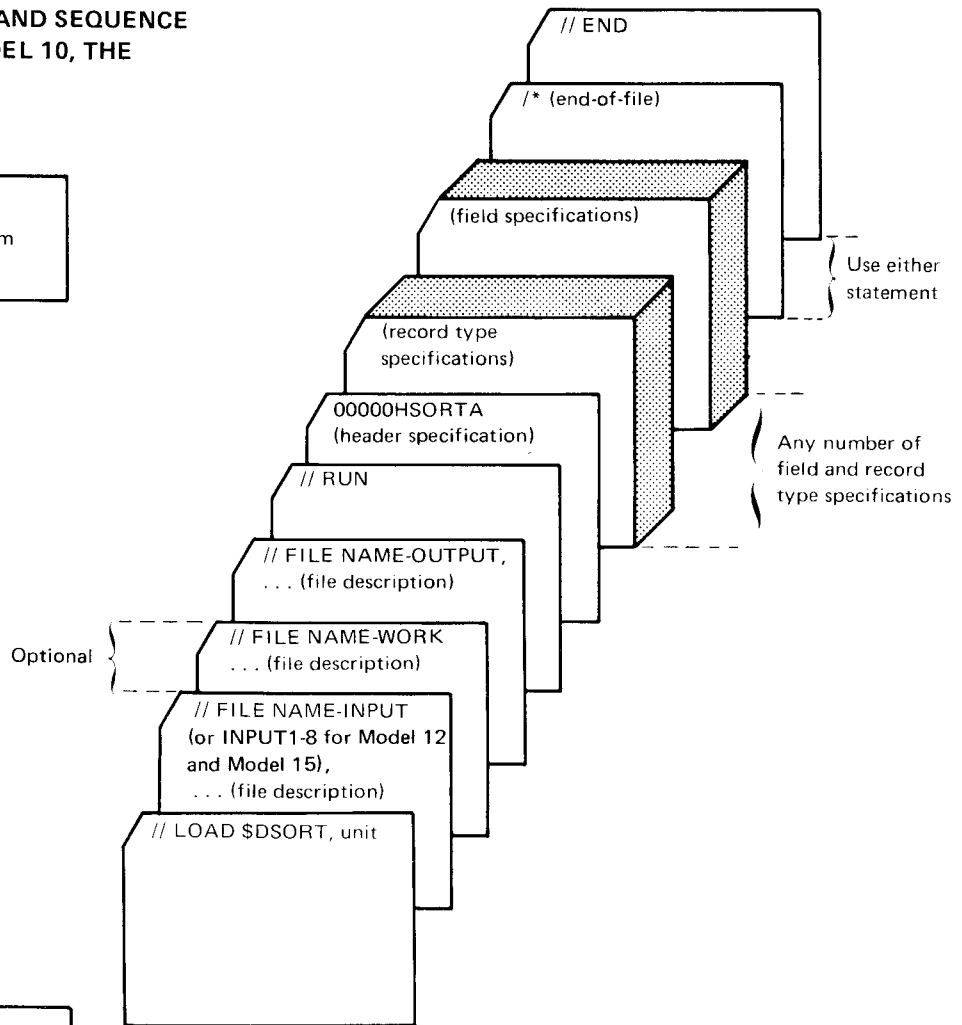
The following chart shows the skeleton of the OCL LOAD and CALL cycles: the first two keywords and any other keywords that are of special significance in the interrelationship of the OCL and the sequence specifications. The rest of the keywords (indicated by arrows) follow the normal sequence for that cycle. (Appendix A contains the keyword-response summaries for each of the three cycles.)

The gray areas in each box show where the sequence specifications fit into the OCL cycle.

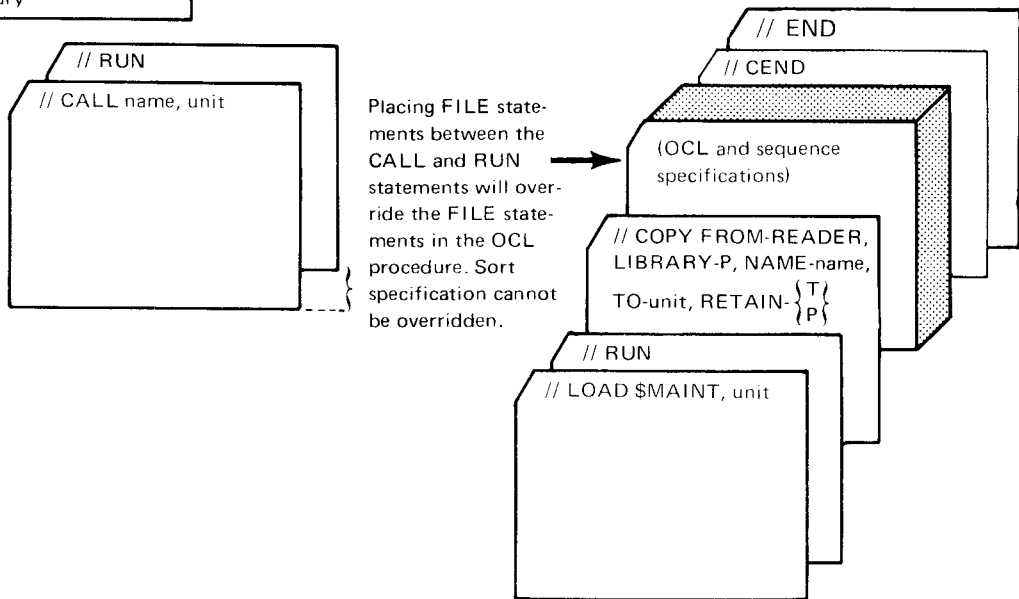


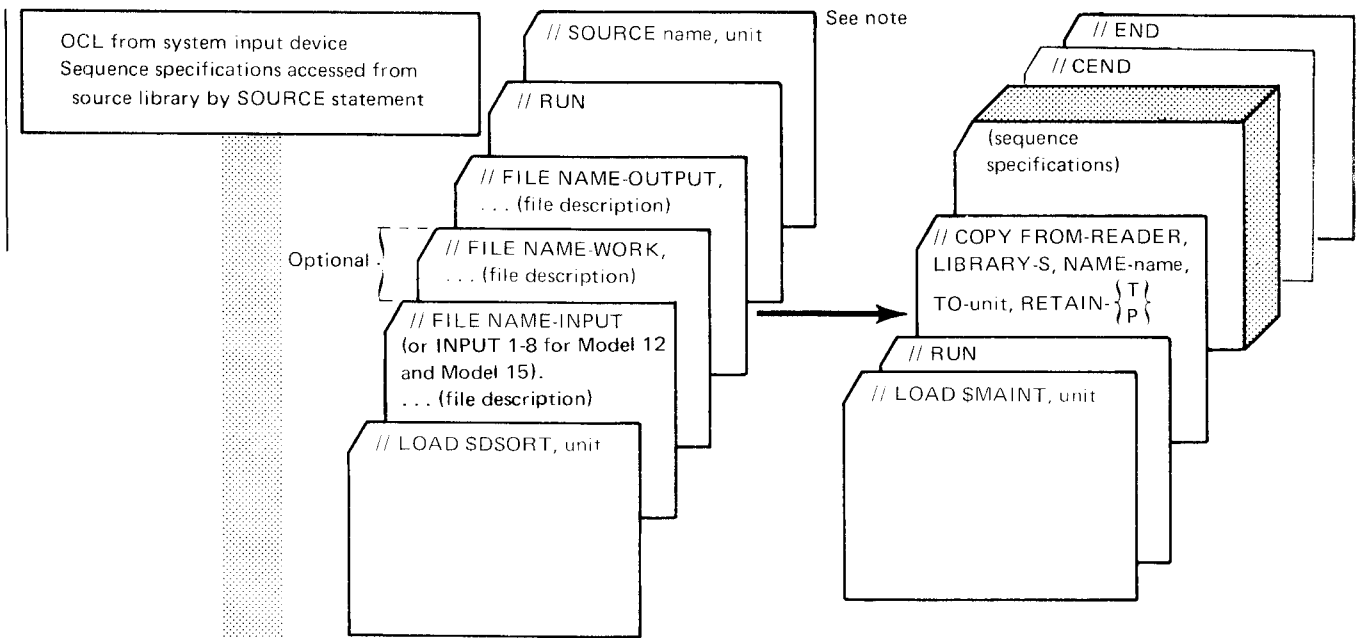
SUPPLYING OCL STATEMENTS AND SEQUENCE SPECIFICATIONS FOR THE MODEL 10, THE MODEL 12, AND THE MODEL 15

OCL from system input device
 Sequence specifications from system input device

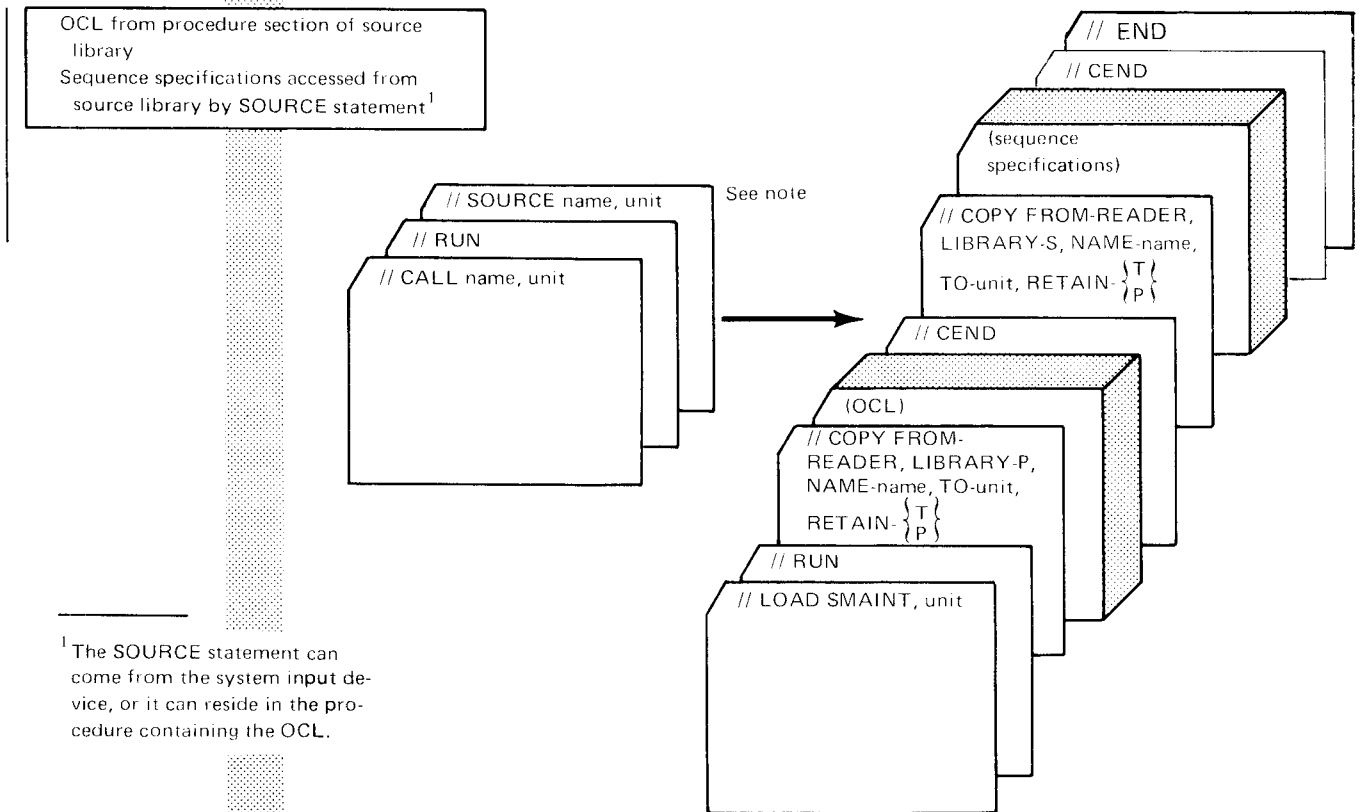


OCL from procedure section of source library
 Sequence specifications from procedure section of source library



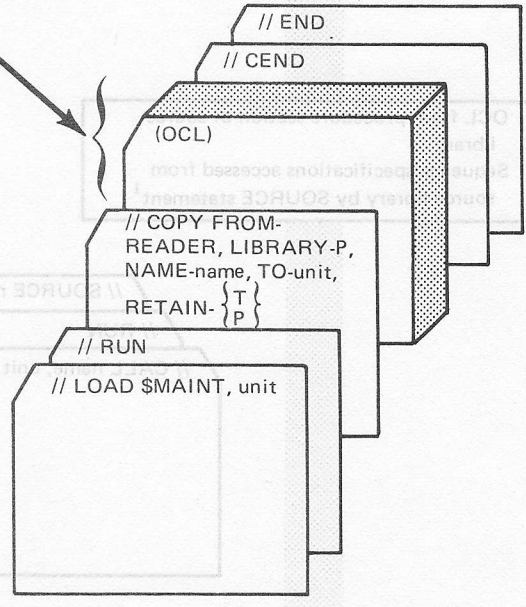
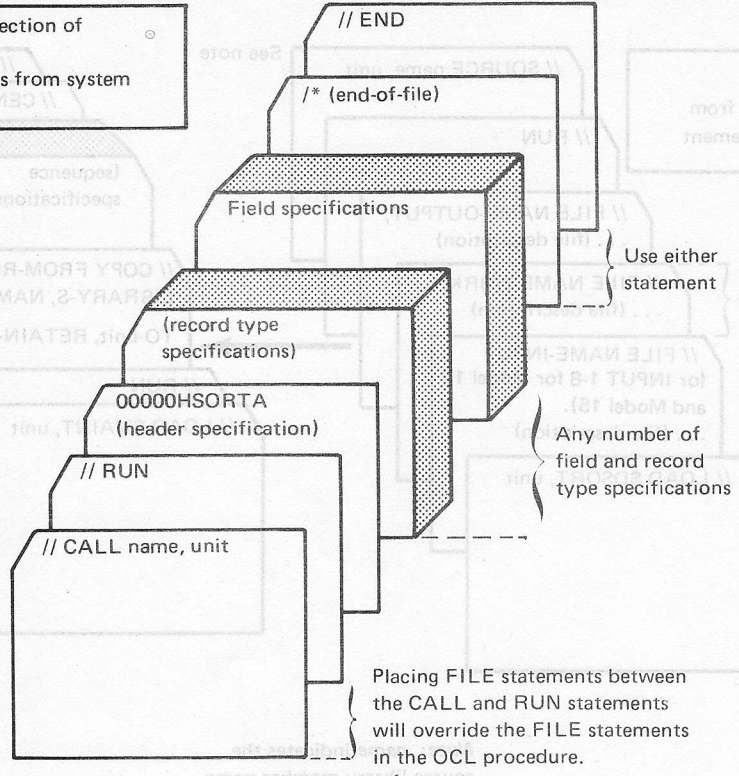


Note: *name* indicates the source library member name containing sort sequence specifications; *unit* indicates the unit identification on which the source library resides.



¹ The SOURCE statement can come from the system input device, or it can reside in the procedure containing the OCL.

OCL from procedure section of source library
Sequence specifications from system input device



WHICH OCL AND SEQUENCE SPECIFICATION METHOD TO USE FOR THE MODEL 10, THE MODEL 12, AND THE MODEL 15

Appendix C. Collating Sequence

The method you use to supply OCL and sequence specifications depends on the type of sort you want to do:

Type of Sort	Method to Use	How the Method Works
Sort jobs you plan to run only a few times	OCL from system input device; sequence specifications from system input device	You supply both OCL and sequence specifications through the system input device.
Sort jobs you plan to run repetitively	Both OCL and sequence specifications from procedure section of source library (maximum of 25 specification statements allowed in procedure)	You place the OCL statements and sequence specifications in the procedure section of the source library using library maintenance, then access them with CALL and RUN statements.
Sort jobs you plan to run using the same file but different sequence specifications	OCL from procedure section of source library; sequence specifications from system input device	You place OCL statements in the procedure section of the source library, then access them using CALL and RUN statements. You supply sequence specifications through the system input device.
	OCL from system input device; sequence specifications accessed from source section of source library by SOURCE statement from the system input device	You supply OCL from the system input device. The SOURCE statement accesses sequence specifications which you have placed in the source section of the source library.
	OCL from the procedure section of source library; sequence specifications accessed from source section of source library by SOURCE statement from the system input device or from the procedure section of the source library	You place OCL statements in the procedure section of the source library, then access them using CALL and RUN statements. The SOURCE statement accesses sequence specifications which you have placed in the source section of the source library.

When several characters share the same position in the sequence, they are considered equal. For example, if you are using only the digit portion of character, H, O, Y and 8 (position 8) are considered equal. This is the number you use in ALTSOQ statements to identify a character you want to shift to a different order in the sequence.

Appendix C. Collating Sequence

Standard EBCDIC Collating Sequence								
When Both Zone and Digit Portions of Characters are Used			When Only the Zone Portion of Characters is Used			When Only the Digit Portion of Characters is Used		
Order in the Sequence ¹	Character	Corresponding Hexadecimal Number ²	Order in the Sequence ¹	Character	Corresponding Hexadecimal Number ²	Order in the Sequence ¹	Character	Corresponding Hexadecimal Number ²
1 (lowest)	blank	40	1 (lowest)	⊘	4A	1 (lowest)	blank	40
2	⊘	4A		&	4B		- (minus)	50
3	.	4B		<	4C)	60
4	<	4C		(4D		{	D0
5	(4D	2	+	4E	2		F0
6	+	4E		!	5A		0	F0
7		4F		*	5B		/	61
8	&	50)	5C		A	C1
9	!	5A	3	;	5D	3	J	D1
10	\$	5B]	5E		1	F1
11	*	5C		/	5F		B	C2
12)	5D		%	61		K	D2
13	;	5E	4	⌵	6B	4	2	F2
14	⌵	5F		⌵	6C		C	C3
15	- (minus)	60		(underscore)	6C		L	D3
16	/	61		>	6E		T	E3
17	,	6B	5	? (apostrophe)	6F	5	3	F3
18	%	6C		#	7A		D	C4
19	⌵ (underscore)	6D		@	7B		M	D4
20	>	6E		'	7C		U	E4
21	?	6F	6	⌵ (apostrophe)	7D	6	4	F4
22	:	7A		⌵	7E		E	C5
23	#	7B		⌵	7F		N	D5
24	@	7C		⌵	7F		V	E5
25	' (apostrophe)	7D	7	⌵	7F	7	5	F5
26	=	7E		⌵	7F		F	C6
27	"	7F		⌵	7F		O	D6
28	A	C1		⌵	7F		W	E6
29	B	C2	8 (highest)	⌵	7F	8	6	F6
30	C	C3		⌵	7F		G	C7
31	D	C4		⌵	7F		P	D7
32	E	C5		⌵	7F		X	E7
33	F	C6	9	⌵	7F	9	7	F7
34	G	C7		⌵	7F		H	C8
35	H	C8		⌵	7F		Q	D8
36	I	C9		⌵	7F		Y	E8
37	J	D0	10	⌵	7F	10	8	F8
38	J	D1		⌵	7F		I	C9
39	K	D2		⌵	7F		R	D9
40	L	D3		⌵	7F		Z	E9
41	M	D4	11	⌵	7F	11	9	F9
42	N	D5		⌵	7F		⌵	4A
43	O	D6		⌵	7F		!	5A
44	P	D7		⌵	7F		:	7A
45	Q	D8	12	⌵	7F	12	.	4B
46	R	D9		⌵	7F		\$	5B
47	S	E2		⌵	7F		,	6B
48	T	E3		⌵	7F		#	7B
49	U	E4	13	⌵	7F	13	<	4C
50	V	E5		⌵	7F		*	5C
51	W	E6		⌵	7F		%	6C
52	X	E7		⌵	7F		@	7C
53	Y	E8	14	⌵	7F	14	(4D
54	Z	E9		⌵	7F)	5D
55	0	F0		⌵	7F		⌵ (underscore)	6D
56	1	F1		⌵	7F		' (apostrophe)	7D
57	2	F2	15	⌵	7F	15	+	4E
58	3	F3		⌵	7F		;	5E
59	4	F4		⌵	7F		>	6E
60	5	F5		⌵	7F		=	7E
61	6	F6	16 (highest)	⌵	7F	16 (highest)		4F
62	7	F7		⌵	7F		⌵	5F
63	8	F8		⌵	7F		?	6F
64 (highest)	9	F9		⌵	7F		"	7F

¹When several characters share the same position in the sequence, they are considered equal. For example, if you are using only the digit portion of characters; H, Q, Y and 8 (position 9) are considered equal.

²This is the number you use in ALTSEQ statements to identify a character you want to shift to a different order in the sequence.

STANDARD COLLATING SEQUENCE

The standard collating sequence chart shows the three standard collating sequences for System/3:

- Standard collating sequence when both zone and digit portions of characters are being used
- Standard collating sequence when only the zone portion of characters is being used
- Standard collating sequence when only the digit portion of characters is being used

ALTERNATE COLLATING SEQUENCE

If you want to change the standard collating sequence, you must code an S in column 26 of the header specification and code ALTSEQ statements. You can code as many ALTSEQ statements as you need. Each one must begin with ALTSEQ and can contain a maximum of 96 positions.

Note: Do not use a packed or unpacked factor 1 in an include or omit record type specification (P or U in column 8) if you specify an alternate collating sequence.

Coding Rules


1. Code ALTSEQ in the first six positions to tell disk sort you want to alter the standard collating sequence:



2. Leave the next two positions blank:



3. Enter the hex equivalent¹ of the character you are taking out of its normal sequence:




4. Enter the hex equivalent¹ of the value the character specified in columns 9 and 10 is to assume in the collating sequence:




5. Enter as many pairs of 3 and 4 as the number of characters you are taking out of normal sequence.

6. Leave no spaces between sets of hex numbers:



7. When you reach the end of one statement, you can continue on the next specification line (follow rules 1-6).

8. Enter a double asterisk in positions 1 and 2 to indicate the end of the ALTSEQ statements:



¹ See standard collating chart for hex equivalents of all System/3 characters.

Input Order When ALTSEQ Statements Are Used

ALTSEQ statements follow the header line. When you use an alternate collating sequence, your input must be in this order:

1. OCL statements
2. Sequence specifications
 - a. Header line
 - b. ALTSEQ statements
 - c. **
 - d. Record type and field description lines as required
3. File to be sorted

Programming Considerations

When you move a character into the sequence position normally assigned to another, both the new and the original character occupy the same position and are considered equal (see the first sample ALTSEQ statement). If you do not want the two characters to be equal, you must also move the character that normally occupies that position (see the second ALTSEQ statement).

Programming Considerations With Card or 3741 Input

Whenever a record is padded on the right with blanks, the padded area is altered by the ALTSEQ logic.

Effect of ALTSEQ Statements on Other Coding

ALTSEQ statements can change:

- Factor 1 and factor 2
- Normal and opposite control fields
- Control field characters before they are replaced or added to by forced fields

ALTSEQ statements never change data fields in records or forced control field characters.

Sample ALTSEQ Statements

To change the sequence of one character:

<i>ALTSEQ</i>	ALTSEQ 40F0
<i>Statements</i>	**

<i>Explanation</i>	Blank (40) is moved to the position in the collating sequence occupied by zero (F0). Blanks and zeros, therefore, are considered equal.
--------------------	---

To change the sequence of several characters:

<i>ALTSEQ</i>	ALTSEQ 5BC8C8C9C9D1
<i>Statements</i>	**

<i>Explanation</i>	5BC8 moves \$ (5B) to the position occupied by H (C8). C8C9 moves H (C8) to the position occupied by I (C9). C9D1 moves I (C9) to the position occupied by J (D1). The \$, therefore, has been inserted between G and H, as follows: A,B,C,D,E,F,G,\$,H,I, (I and J occupy the same position).
--------------------	--

PROGRAMMING AIDS

A Translation Table and Alternate Collating Sequence Coding Sheet, X21-9096, is available for you to use to keep track of changes you make to the collating sequence.

TRANSLATION TABLE AND ALTERNATE COLLATING SEQUENCE CODING SHEET

Code	Graphic	Entry	Replaced By/Takes Place Of
00000000		00	
00000001		01	
00000010		02	
00000011		03	
00000100		04	
00000101		05	
00000110		06	
00000111		07	
00001000		08	
00001001		09	
00001010		0A	
00001011		0B	
00001100		0C	
00001101		0D	
00001110		0E	
00001111		0F	
00010000		10	
00010001		11	
00010010		12	
00010011		13	
00010100		14	
00010101		15	
00010110		16	
00010111		17	
00011000		18	
00011001		19	
00011010		1A	
00011011		1B	
00011100		1C	
00011101		1D	
00011110		1E	
00011111		1F	
00100000		20	
00100001		21	
00100010		22	
00100011		23	
00100100		24	
00100101		25	
00100110		26	
00100111		27	
00101000		28	
00101001		29	
00101010		2A	
00101011		2B	
00101100		2C	
00101101		2D	
00101110		2E	
00101111		2F	
00110000		30	
00110001		31	
00110010		32	

Code	Graphic	Entry	Replaced By/Takes Place Of
01100010		55	
01100011		56	
01101000		57	
01101001		58	
01101010		59	
01101011		6A	
01101100		6B	
01101101		6C	
01101110		6D	
01101111		6E	
01110000		6F	
01110001		70	
01110010		71	
01110011		72	
01110100		73	
01110101		74	
01110110		75	
01110111		76	
01111000		77	
01111001		78	
01111010		79	
01111011		7A	
01111100		7B	
01111101		7C	
01111110		7D	
01111111		7E	
10000000		80	
10000001		81	
10000010		82	
10000011		83	
10000100		84	
10000101		85	
10000110		86	
10000111		87	
10001000		88	
10001001		89	
10001010		8A	
10001011		8B	
10001100		8C	
10001101		8D	
10001110		8E	
10001111		8F	
10010000		90	
10010001		91	
10010010		92	
10010011		93	
10010100		94	
10010101		95	
10010110		96	
10010111		97	
10011000		98	

Code	Graphic	Entry	Replaced By/Takes Place Of
10011001		99	
10011010		9A	
10011011		9B	
10011100		9C	
10011101		9D	
10011110		9E	
10011111		9F	
10100000		A0	
10100001		A1	
10100010		A2	
10100011		A3	
10100100		A4	
10100101		A5	
10100110		A6	
10100111		A7	
10101000		A8	
10101001		A9	
10101010		AA	
10101011		AB	
10101100		AC	
10101101		AD	
10101110		AE	
10101111		AF	
10110000		B0	
10110001		B1	
10110010		B2	
10110011		B3	
10110100		B4	
10110101		B5	
10110110		B6	
10110111		B7	
10111000		B8	
10111001		B9	
10111010		BA	
10111011		BB	
10111100		BC	
10111101		BD	
10111110		BE	
10111111		BF	
11000000		C0	
11000001		C1	
11000010		C2	
11000011		C3	
11000100		C4	
11000101		C5	
11000110		C6	
11000111		C7	
11001000		C8	
11001001		C9	
11001010		CA	
11001011		CB	

Code	Graphic	Entry	Replaced By/Takes Place Of
11001100		CC	
11001101		CD	
11001110		CE	
11001111		CF	
11010000		D0	
11010001		D1	
11010010		D2	
11010011		D3	
11010100		D4	
11010101		D5	
11010110		D6	
11010111		D7	
11011000		D8	
11011001		D9	
11011010		DA	
11011011		DB	
11011100		DC	
11011101		DD	
11011110		DE	
11011111		DF	
11100000		E0	
11100001		E1	
11100010		E2	
11100011		E3	
11100100		E4	
11100101		E5	
11100110		E6	
11100111		E7	
11101000		E8	
11101001		E9	
11101010		EA	
11101011		EB	
11101100		EC	
11101101		ED	
11101110		EE	
11101111		EF	
11110000		F0	
11110001		F1	
11110010		F2	
11110011		F3	
11110100		F4	
11110101		F5	
11110110		F6	
11110111		F7	
11111000		F8	
11111001		F9	
11111010		FA	
11111011		FB	
11111100		FC	
11111101		FD	
11111110		FE	
11111111		FF	

Note: Not all graphic symbols shown here are available on all systems.



Hexadecimal numbers of characters in the standard collating sequences.
Where you record the hexadecimal number of the character you are going to put in that relative position in the sequence.

Appendix D. Timing Considerations

The time it takes to run a sort job can vary greatly and is dependent on the following factors:

- The amount of main storage available for the sort program
- Number of records to be sorted
- Record size
- Number of sequence specifications
- Whether alternate collating sequence is being used
- Disk location of the input, work and output files
- Order of control fields in the input file
- Whether or not automatic work file allocation is being used
- Use of nonverify option
- Number of input files and the types of input devices (Model 12 and Model 15)

Here is how each factor affects the length of a sort job.

MAIN STORAGE SIZE

The more main storage the program has:

- The more sequence specifications you can have
- The longer your records can be
- The faster your sort job will run

See Appendix E for more information and detailed charts.

NUMBER OF RECORDS TO BE SORTED

The more records you want to sort, the longer the sort job will take. For the most efficient sort job, therefore, use the record omit option whenever possible.

RECORD SIZE

The larger the records, the longer it will take to sort a file. So, if program run time is important to you, do not include fields that contain unnecessary information on your field specifications. (Every field you do not include decreases the record size by the length of that field.) Remember, disk sort builds a work record for each input record included in the sort.

NUMBER OF SEQUENCE SPECIFICATIONS

The more sequence specifications you use, the less main storage space will be available for records.

- Each sequence specification uses about 25 bytes of storage (remember that comment lines use no main storage).
- The select/build routine uses approximately 100 bytes of storage for functional interfacing.
- If you have specified includes and/or omits (referring to packed data) for your sort run, the select/build routine will use an additional 200 bytes of storage.
- If you have specified field statements (referring to packed data) for your sort run, the select/build routine will use an additional 80 bytes of storage.
- Any alternate collating sequence uses an additional 375 bytes of storage.

But consider using additional sequence specifications to exclude records you do not want to sort. The result may well be a shorter program run than if you use fewer specifications but process more records.

ALTERNATE COLLATING SEQUENCE

Using an alternate collating sequence increases the time it takes to run a sort job.

DISK LOCATION OF THE FILES

Because the disk sort program must move each record you want to sort two times (first from the input file to the work file and then from the work file to the output file), file location is an important factor in how long it takes to run a job. It is important to minimize the distance the disk arm must travel between the input and work file and between the work and output file. The following diagrams show different ways of locating the input, work, and output files for a sort job.


Within each diagram, the different methods of file placement are shown in order of usual performance characteristics. Each succeeding method will normally run faster than the methods preceding it. For systems having one (or two) 5445 disk storage units available for use with sort, work files should be placed on a 5445 disk storage unit, if possible. If D1 and D2 are both available, put the work file on one drive and the input and output files on the other.

Amount of Main Storage Assigned to Disk Sort ¹	Maximum Recommended Number of Bytes of Select/Build Code
5K	768
6K	768
7K	1024
8K	1280
9K	1536
10K	1792
11K	2048
12K	2304
13K	2560
14K	2816
15K	3072
16K	3328
17K	3584
18K	3840
19K and above	4096


¹ Maximum storage in the Model 6 is 16K

5444 File Placement


Using one disk on a single drive system




Using two disks on a single drive system



Using three disks on a dual drive system




Using four disks on a dual drive system




5445/3340 File Placement

Using one drive



Using two drives



Note: If the input and output files are on drives separate from the work file, this is optimal file placement. If additional 5445 or 3340 disk drives are available (up to four on Model 15), then much larger files can be sorted and optimal file placement is easier.

ORDER OF THE CONTROL FIELDS IN THE INPUT FILE

Suppose two files contain identical information; however, some of the first file's records and fewer of the second file's records are in the sequence you want. *Using the same sequence specifications* to sort each file, disk sort takes less time to sort the first file.

USE OF THE AUTOMATIC WORK FILE ALLOCATION FUNCTION

Using the automatic work file allocation function generally increases the time needed to run a sort job, because this function does not always provide the work file arrangement needed for a fast sort run. If you are concerned with minimizing sort run-time, use a well planned work file and work FILE statement, rather than automatic work file allocation.

NONVERIFY OPTION

If an N is entered in column 34 of the disk sort header statement, the disk sort program will not verify data written on the work file. The nonverify option can be used on either a 5444 or 5445 disk pack. The amount of performance improvement depends on file size, number of records, file locations, etc.

Note that when a work file statement is specified for a 3340 file (D1, D2, etc) on Model 15, the verify option is taken from the file statement and column 34 is ignored.

SAMPLE DISK SORT TIMINGS

The following timings are supplied to aid you in approximating job run times. The sort times shown were generated under the following conditions; any changes may decrease performance:

- Core sizes shown refer to the amount of main storage used by disk sort, excluding the supervisor.
- All timings reflect the total sort run time (generation time plus execution time).

- All runs were made using SORTR, and the record lengths shown reflect input, work, and output record lengths.
- All control fields used were ten characters long.
- All input records were included for minimum select/build generation.
- Alternate collating sequence was not used.
- No alternate track assignments were made.
- Automatic work file allocation was not used.
- Optimal file placement was used for all runs (see *Disk Location of the Files*).

Model 6 and Model 10 Timings Only

- Input, work, and output files resided on 5444 Model 2 and 5445 disks for 5444 and 5445 timings, respectively.
- None of the 5445 disks are split cylinder files.
- The input files was overlaid by the output file to minimize file space usage.
- All files were sorted into inverse sequence (from ascending to descending sequence).

Model 15 Timings

- Input and output files were on D1; work files were on D2.
- Only one partition was active during the timings.
- All OCL and sort specifications were put on a spool file on D3.
- All files were sorted into descending sequence (from random sequence to descending sequence).

Sample Model 6 and Model 10 Timings

		Record Size			
		Single Drive		Dual Drive	
		96 bytes	200 bytes	96 bytes	200 bytes
5444	5K Partition				
	5,000 Records	18 min.	49 min.	11 min.	30 min.
	10,000 Records	34	107	21	62
	12K Partition				
	5,000 Records	7	15	7	11
	10,000 Records	14	31	11	21
5445	28K Partition				
	5,000 Records	5	8	5	8
	10,000 Records	9	17	9	13
	12K Partition				
	10,000 Records	9	14	9	12
	30,000 Records	32	52	25	33
5445	28K Partition				
	10,000 Records	7	10	7	10
	30,000 Records	17	31	17	26
	44K Partition				
	10,000 Records	7	9	6	9
	30,000 Records	16	26	15	24

Note: All times are rounded to the nearest minute.

Sample Model 15 Timings

		Record Size					
		100 bytes		200 bytes			
Model 15A 5445 Timing	24K Partition	5,000 Records	3 Min.	38s	4 Min.	56s	
		10,000 Records	5	24	8	28	
		20,000 Records	10	1	17	12	
		50,000 Records	23	23	-	-	
	48K Partition	5,000 Records	3	4	3	47	
		10,000 Records	4	28	6	48	
		20,000 Records	7	50	11	37	
		50,000 Records	19	10	31	25	
	Model 15B and C 3340 Timing	24K Partition	5,000 Records	1	52	2	31
			10,000 Records	3	9	5	18
			20,000 Records	6	17	9	40
			50,000 Records	15	53	-	-
48K Partition		5,000 Records	1	31	1	56	
		10,000 Records	2	36	3	43	
		20,000 Records	5	3	7	1	
		50,000 Records	12	42	18	55	
Model 15D 3340 Timing		24K Partition	5,000 Records	1	34	2	19
			10,000 Records	2	43	4	13
			20,000 Records	5	9	8	42
			50,000 Records	13	26	23	14
	48K Partition	5,000 Records	1	21	1	48	
		10,000 Records	2	16	3	17	
		20,000 Records	4	16	6	6	
		50,000 Records	10	46	16	30	

Note: These sample timings increase in direct proportion to the number of records. For example, using the 5444 (12K partition with a single drive) to sort 7500 96-byte records would take approximately 10 minutes. That is, about 50% more records are being sorted (50% of 5000 = 2500; 2500 + 5000 = 7500) so the sorting time takes 50% longer (50% of 7 = about 3; 3 + 7 = 10).

Sample Model 12 Timings

Use the sample timings illustrated for Model 15B for a reasonable approximation for the Model 12 run times.

Appendix E. Calculating Size of Files

INPUT FILE

Use this formula to calculate how many tracks the input file will require:

$$\frac{\text{Number of records in the file} \times \text{input record length}}{\text{Number of bytes in a track}}$$

The number of bytes in a track is 6144 for the 5444, 5120 for the 5445, and 12,288 for the 3340 (Model 12 and Model 15). Round the result up to the nearest whole number.

OUTPUT FILE

Use this formula to calculate how many tracks the output file will require:

$$\frac{\text{Number of records to be sorted} \times \text{output record length}}{\text{Number of bytes in a track}}$$

The number of bytes in a track is 6144 for the 5444, 5120 for the 5445, and 12,288 for the 3340 (Model 12 and Model 15). Round the result up to the nearest whole number.

Note: Number of records to be sorted means the number of records that will be selected to be sorted.

WORK FILE

Do not use RECORDS to specify the work file, use TRACKS.

Use this formula to approximate how many tracks the work file will require:

$$2 + \left(\frac{\text{Number of records to be sorted} \times \text{work record length} \times \text{approximate work file size factor}}{\text{Number of bytes in a track}} \right)$$

The number of bytes in a track is 6144 for the 5444, 5120 for the 5445, and 12,288 for the 3340 (Model 12 and Model 15). Round the result up to the nearest whole number.

Rather than using this formula to calculate the number of work file tracks needed for each additional run of this job, you can use the actual number of tracks of work file used in this sort. This number is automatically calculated at the end of first pass #0 and is printed out as message SD695. In order to save time, you can use automatic work file allocation for the first run. Then, in subsequent runs, use work FILE statements specifying the number of tracks printed in message SD695.

Work Record Length

Type of Sort	Type of Output	Work Record Length
SORTA	Address output	Length of control fields +3
SORTR or SORTRS	Control fields only	Length of control fields as specified in the header line
	Control fields and data	Length of data + length of control fields. (This is the output record length specified in the header line.)
	Data only	Length of data + length of control fields. (Output record length specified in header line plus the control field length specified in header line.)

The following table shows the maximum work record length for a disk sort run, using various combinations of main storage size, input record length, and number of bytes of select/build code. For example, if disk sort is assigned 12K of main storage and your input record length is 1600 bytes, the maximum work record length you can have, using the recommended maximum number of bytes of select/build code, is 2808 bytes. However, if you used

fewer than 256 bytes of select/build code in the same situation, you can have a work record length of 3064 bytes. Generally, then, this table shows that for a given amount of main storage, the maximum work record length can be increased by using either fewer bytes of select/build code, or smaller input records, or both. This table can also be used to estimate the maximum input and output block size for tape input/output.

Amount of Main Storage Assigned to Disk Sort	INPUT RECORD LENGTH		Maximum Work Record Length
	For Jobs with Fewer than 256 Bytes of Select/Build Code	For Jobs with Maximum Recommended Number of Bytes of Select/Build Code ¹	
6K	01-768	01-384	760 ²
	769-896	385-640	504
	897-1024	641-768	248 ³
8K	01-1792	01-896	1784 ²
	1793-1920	897-1408	1528
	1921-2048	1409-1536	1272
	2049-2176	1537-1664	1016
	2177-2304	1665-1792	760
	2305-2432	1793-1920	504
9K	2433-2560	1921-2048	248 ³
	01-2176	01-1152	2040
	2177-2304	1153-1664	1784
	2305-2432	1665-1792	1528
	2433-2560	1793-1920	1272
	2561-2688	1921-2048	1016
	2689-2816	2049-2176	760
12K	2817-2944	2177-2304	504
	2945-3072	2305-2432	248 ³
	01-3200	01-1408	3064
	3201-3328	1409-1920	2808
	3329-3456	1921-2432	2552
	3457-3584	2433-2560	2296
	3585-3712	2561-2688	2040
	3713-3840	2689-2816	1784
	3841-3968	2817-2944	1528
	3969-4096	2945-3072	1272
20K and up		3073-3200	1016
		3201-3328	760
		3329-3456	504
		3457-3584	248 ³
	01-4096	01-4096	4096

¹ See *Number of Sequence Specifications* in Appendix D for maximum recommended number of bytes of select/build code per job.

² Not valid for SORTRS (summary sort).

³ For SORTA (addroot sort), 245 is the largest total length of control fields for any record type.

Work File Size Factor

Amount of Main Storage Assigned to Disk Sort	INPUT RECORD LENGTH		Approximate Work File Size Factor
	For Jobs with Less than 11 Sequence Specifications	For Jobs with 11 or More Sequence Specifications	
6K	01-768	01-384	1.52
	769-1024	385-768	2.05
8K	01-1792	01-896	1.17
	1793-1920	897-1408	1.20
	1921-2048	1409-1536	1.25
	2049-2176	1537-1664	1.35
	2177-2304	1665-1792	1.52
	2305-2432	1793-1920	2.05
	2433-2560	1921-2048	2.05
9K	01-2176	01-1152	1.15
	2177-2304	1153-1664	1.17
	2305-2432	1665-1792	1.20
	2433-2560	1793-1920	1.25
	2561-2688	1921-2048	1.35
	2689-2816	2049-2176	1.52
	2817-2944	2177-2304	2.05
2945-3072	2305-2432	2.05	
12K	01-3200	01-1408	1.09
	3201-3328	1409-1920	1.10
	3329-3456	1921-2432	1.12
	3457-3584	2433-2560	1.13
	3585-3712	2561-2688	1.15
	3713-3840	2689-2816	1.17
	3841-3968	2817-2944	1.20
	3969-4096	2945-3072	1.25
		3073-3200	1.35
		3201-3328	1.52
	3329-3456	2.05	
	3457-3584	2.05	
20K and up	01-4096	01-4096	1.07

EXAMPLE

For this example, assume these values:

Number of records in the file or files (Model 15 only) = 2000

Number of records to be sorted = 1000

Input record length = 200 bytes

Work record length:

SORTA = 13 bytes (control field length = 10)
 SORTR = 210 bytes (control field length = 10)

Output record length:

SORTA = 3 bytes
 SORTR = 200 bytes (drop control field)

Main storage size = 8K

Work file size factor = 1.17

Here are the calculations:

5444 Disk Files

c

	SORTR	SORTA
Input file	$\left(\frac{2000 \times 200}{6144} \right) = 65.1$ 66 tracks	$\left(\frac{2000 \times 200}{6144} \right) = 65.1$ 66 tracks
Work file	$\left(\frac{1000 \times 210}{6144} \right) \times (1.17) + 2 = 41.99$ 42 tracks	$\left(\frac{1000 \times 13}{6144} \right) \times (1.17) + 2 = 4.48$ 5 tracks
Output file	$\left(\frac{1000 \times 200}{6144} \right) = 32.6$ 33 tracks	$\left(\frac{1000 \times 3}{6144} \right) = 0.49$ 1 track
	Total tracks for all files = 141	Total tracks for all files = 72

5445 Disk Files

	SORTR	SORTA
Input file	$\left(\frac{2000 \times 200}{5120}\right) = 78.1$ 79 tracks	$\left(\frac{2000 \times 200}{5120}\right) = 78.1$ 79 tracks
Work file	$\left(\frac{1000 \times 210}{5120}\right) \times (1.17) + 2 = 49.99$ 50 tracks	$\left(\frac{1000 \times 13}{5120}\right) \times (1.17) + 2 = 4.97$ 5 tracks
Output file	$\left(\frac{1000 \times 200}{5120}\right) = 39.1$ 40 tracks	$\left(\frac{1000 \times 3}{5120}\right) = .59$ 1 track
	Total tracks for all files = 169	Total tracks for all files = 85

3340 Disk Files (Model 12 and Model 15)


	SORTR	SORTA
Input file	$\left(\frac{2000 \times 200}{12288}\right) = 32.6$ 33 tracks	$\left(\frac{2000 \times 200}{12288}\right) = 32.6$ 33 tracks
Work file	$\left(\frac{1000 \times 210}{12288}\right) \times (1.17) + 2 = 19.1$ 20 tracks	$\left(\frac{1000 \times 13}{12288}\right) \times (1.17) + 2 = 3.3$ 4 tracks
Output file	$\left(\frac{1000 \times 200}{12288}\right) = 16.3$ 17 tracks	$\left(\frac{1000 \times 3}{12288}\right) = .3$ 1 track
	Total tracks for all files = 70	Total tracks for all files = 38

An alphabetic code helps you judge the significance of each disk sort message:

Code	Meaning	Significance
T	Terminal	The error was destructive. The program must end immediately. No more error messages will be printed.
S	Severe Error	A source statement contains an error. The program ends after the generation phase.
A	Action	The operator must do something. A system halt follows this message.
W	Warning	An unusual condition has been found. It should be checked out to make sure it did not occur as a result of a programming error.
I	Information	For information only. Messages with an I code indicate current status and progress of job.

To understand the messages, you must also know something about how the disk sort program works. Specifically, you must know what happens during the generation phase of the program and what happens during the execution phase.

The generation phase translates your sequence specifications into machine language:



The execution phase gets the records you want sorted, sorts the records, and finally writes the sorted records onto the output file in the order you want them.

¹The object program is often referred to as the select/build routine.

**SD100—ENTER SORT SEQUENCE SPECIFICATIONS—
HEADER FOLLOWED BY INCLUDE/OMIT SETS OR
// SOURCE NAME, UNIT STATEMENT**

Code: A—Action
Explanation: You must tell the system where the sequence specifications for the job are. There are two possibilities: the operator is going to type them on the keyboard (Model 6), on the printer/keyboard (Model 10 or Model 12), or they are already in a source library on one of the disks.
System Action: System waits for response from operator.
User Response: If the operator is going to type the specifications on the keyboard (Model 6), or printer/keyboard (Model 10 or Model 12), he or she must follow this sequence:
1. Header line
2. All the include/omit sets
3. // END
If the sequence specifications are already in a source library on one of the disks the operator types a SOURCE statement. This gives the system the name and disk drive location of the set of sequence specifications.

SD101—PRINT OPTION ASSUMED TO BE 0

Code: I—Information
Explanation: The print option in your header line is not blank 0, 1, 2 or 3; therefore, it is assumed to be 0.
System Action: The job continues.
User Response: None required.

SD102—SORT JOB ASSUMED TO BE SORTR

Code: W—Warning
Explanation: Neither SORTA, SORTR, or SORTRS job type is specified on your header line. Because an output record length is specified, SORTR is assumed.
System Action: The job continues.
User Response: Check to make sure that a SORTR job was wanted. To avoid this error the next time the job is run, fill in columns 7-12 of the header line properly.

SD103—SORT JOB TYPE NOT SPECIFIED

Code: T—Terminal
Explanation: Neither a job type nor an output record length is specified in the header line.
System Action: The job ends.
User Response: Place the proper job type in columns 7-12 of your header line. Resubmit job.

**SD104—BOTH INPUT AND INPUT1 ARE PRESENT
(Model 12 and Model 15)**

Code: S—Severe Error
Explanation: The filenames INPUT and INPUT1 are considered equivalent by disk sort. Only one of these filenames may be used in a disk sort.
System Action: The job ends after the generation phase.
User Response: Eliminate one of the FILE statements or change one of the filenames to INPUT2, INPUT3, . . . , INPUT8.

SD105—INPUT FILE NOT ONLINE(Disk Sort only)

Code: T—Terminal
Explanation: Pack (or packs) containing the input file is not mounted.
System Action: Halt 25 occurs for Model 10, Model 12, or Model 15, halt CD45 for Model 6.
User Response: Mount the pack (or packs) containing the input file. If the input file is multi-volume offline (if more than one pack is assigned to a removable unit), only the first pack assigned to each unit must be mounted. Resubmit the job.

SD106—WORK FILE NOT ON DISK (Not CCP/Disk Sort)

Code: T—Terminal
Explanation: The work FILE statement specifies that the work file is to reside on a device other than disk, and disk sort requires the work file to be on disk.
System Action: The job ends.
User Response: Either omit the work FILE statement or specify the work file to be on disk. If you omit the work FILE statement, disk sort automatically allocates your work file.

SD108—INPUT AND/OR OUTPUT FILE STATEMENT MISSING

Code: T—Terminal
Explanation: Either the input or output (or both) FILE statement is missing in this run.
System Action: The job ends. The CD45 halt (Model 6) or 25 halt (Model 10, Model 12, or Model 15) cancels the job.
User Response: Include missing FILE statement(s) and resubmit the job.

SD109—WORK FILE STATEMENT MISSING

Code: S—Severe error
Explanation: A work file was not specified.
System Action: The job ends after the generation phase.
User Response: Include the missing FILE statement and resubmit the job.

SD150—CONTROL FIELD DROPPED—NO DATA SPECIFICATIONS

Code: S—Severe
Explanation: No data specifications (summary specifications included) were given, yet the control field is being dropped. If the sort were allowed to continue, the output file would be all blanks (X'40').
System Action: The job ends after the generation phase.
User Response: Add data field specifications and/or retain the control field; then resubmit the job.

SD151—CONTROL FIELD LENGTH NOT SPECIFIED

Code: T—Terminal
Explanation: No control field length specified on header line.
System Action: The job ends.
User Response: You must place the length of your control fields in columns 13-17. Do so, and re-submit job.

SD152—CONTROL FIELD LENGTH TOO LARGE

Code: S—Severe error
Explanation: Control field specified too large; 256 is assumed.
System Action: The job ends after the generation phase.
User Response: Make the control field length in columns 13-17 256 or less. Resubmit job.

SD153—INPUT RECORD LENGTH UNDETERMINABLE

Code: T—Terminal
Explanation: The record length of your input file cannot be determined from information stored on disk.
If you specified RECORDS or TRACKS on the File statement for the input file and the file was not on the specified pack, sort was processing a newly created file and not your input file.
CCP/Disk Sort users — Columns 39 and 40-43 of the header statement did not contain an R or C followed by the input record length.
System Action: The job ends.
User Response: This is a system-program error. If error persists, contact your IBM representative for program assistance.

**SD154—ASCENDING SEQUENCE ASSUMED A,
COLUMN 18**

Code: W—Warning
Explanation: Column 18 of your header line has neither a D for descending, nor an A for ascending sequence. Disk sort assumes you want ascending sequence.
System Action: The job continues.
User Response: If you want descending sequence, you can cancel the job at the C123 halt (Model 6), or at the 22 halt (Model 10, Model 12, or Model 15).

SD155—OUTPUT OPTION ASSUMED X, COLUMN 28

Code: W—Warning
Explanation: Column 28 of your header line is neither an X nor blank.
System Action: Disk sort assumes you want an X in column 28 and continues the job.
User Response: If you do not want X, you can cancel the job at the C123 halt (Model 6), or at the 22 halt (Model 10, Model 12, or Model 15).

SD156—OUTPUT RECORD LENGTH NOT SPECIFIED

Code: S—Severe error
Explanation: This is a SORTR or SORTRS job. Output record length not specified on header line. Disk sort assumes you want the maximum length of 4096.
System Action: The job ends after the generation phase.
User Response: Enter an output record length in columns 29-32 of your header line and resubmit the job.

SD157—OUTPUT RECORD LENGTH TOO LARGE

Code: S—Severe error
Explanation: This is a SORTR or SORTRS job. Output record length specified on header line is too large. Disk sort assumes you want the maximum length of 4096.
System Action: The job ends after the generation phase.
User Response: Make sure your header line is not shifted by one or more columns. Specify on output record length of 4096 or less and re-submit job.

SD158—OUTPUT RECORD LENGTH INCONSISTENT

Code: S—Severe error
Explanation: This is a SORTR or SORTRS job and you are not dropping the control field. Consequently, the output record length (Columns 29-32) that includes the control field should not be less than the control field (columns 13-17).
System Action: A length equal to that in columns 29-32 plus the largest total control field (columns 13-17) is assumed.
User Response: Recompute your output record length, change it on your header line, and resubmit the job.

SD159—WORK RECORD LENGTH TOO LARGE

Code: S—Severe error
Explanation: Disk sort's work record exceeds the maximum length allowed.
System Action: Halt CD45 (Model 6) or halt 25 (Model 10, Model 12, or Model 15) occurs.
User Response: Refer to Appendix E to determine the maximum work record length.

SD160—ALTERNATE COLLATING SEQUENCE ASSUMED

Code: W—Warning
Explanation: Columns 26 is neither an S nor blank.
System Action: Disk sort assumes you want an alternate collating sequence and continues the job.
User Response: If you do want an alternate collating sequence, put an S in column 26 of your header line before running the job again. If you do not want an alternate collating sequence, you can cancel the job at the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15).

SD161—INVALID NUMBER IN CONTROL FIELD

Code: T—Terminal
Explanation: There is either a nonnumeric character or embedded blank in columns 13-17 of the header line.
System Action: The job ends.
User Response: Correct the mistake and resubmit the job.

SD162—INVALID NUMBER IN OUTPUT RECORD LENGTH

Code: T—Terminal
Explanation: There is either a nonnumeric character or embedded blank in columns 29-32 of the header line.
System Action: The job ends.
User Response: Correct the mistake and resubmit the job.

SD163—SORTA-TAPE OUTPUT, RECORD LENGTH NOT 18

Code: T—Terminal
Explanation: The record length on the output FILE statement is not 18.
System Action: The job ends.
User Response: For a tape output file during a SORTA (addrout) sort, the record length on the output FILE statement must be 18.

SD164—SORTA—MULTIPLE INPUTS (Model 12 and CCP/Disk Sort)

Code: S—Severe error
Explanation: Addrout sort is invalid with multiple disk or tape files.
System Action: The job ends after the generation phase.
User Response: Modify your OCL to perform an addrout sort on a single disk or tape file, or merge all the files into a single file before performing an addrout sort.

SD164—SORTA—CARD, 3741, OR MULTIPLE INPUT (Model 15 only)

Code: S—Severe error
Explanation: Addrout sort is invalid with card input, 3741 input, or multiple disk or tape files.
System Action: The job ends after the generation phase.
User Response: Modify your OCL to perform an addrout sort on a single disk or tape file, or merge all the files into a single file before performing an addrout sort.

SD165—SORTA-ASCII TRANSLATE SPECIFIED

Code: T—Terminal
Explanation: If the addrout file is converted to ASCII, the result is an unusable file.
System Action: The job ends.
User Response: ASCII translate must not be specified if a tape addrout file is desired.

SD166--VERIFY OPTION ASSUMED BLANK, COLUMN 34

Code: W--Warning
Explanation: Column 34 of the header does not contain an N or a blank. Disk sort will continue the sort, and assume that you meant for column 34 to be blank. Therefore, any writes to the work file will be verified.
System Action: Disk sort assumes that you want any writes to the work file to be verified.
User Response: If you desire not to have all writes to the work file verified, you must code an N in column 34 and rerun the job.

SD167--INVALID STORAGE SIZE TO EXECUTE

Code: W--Warning
Explanation: The entry in columns 37-38 was:
1. Blank or nonnumeric
2. Less than 12
3. Greater than 48
4. Not a multiple of 2
System Action: For the above cases, the following defaults occur:
-- For items 1 and 2, 12 is assumed.
-- For item 3, 48 is assumed.
-- For item 4, the value is rounded down to the next lowest multiple of 2.
User Response: If the default is not acceptable, enter the proper value and resubmit the job.

SD168--MISSING OR INVALID PROGRAM NAME

Code: W--Warning
Explanation: The entry in columns 75-80 of the header statement is either blank or has an invalid program name.
System Action: The program name defaults to SORTDISJ.
User Response: If the default is not acceptable, enter the correct program name in columns 75-80 and resubmit the job.

SD169--MIXED EBCDIC AND ASCII FILES (Model 10 only)

Code: W--Warning
Explanation: The input file is mapping to mix ASCII and EBCDIC. The input and output files. Both input and output files must be the same type.
System Action: The job will run through the generation phase. At the end of the sort sequence, the job will terminate at the end of the generation phase.
User Response: If you want to use so they are the same type, you must use the same type (EBCDIC). For mixed input and output, translate the input to ASCII. Then run the sort job.

SD172--H-BACK SORT IS SORTA TYPE SORT

Code: W--Warning
Explanation: The input file is SORTA (addroot) sort is not supported. The record number. Binary sort is not supported. HBTZ sort cannot go on.
System Action: The job ends after the generation phase.
User Response: If you want to use the desired, output file, you must use a disk or 9-track tape.

SD173--CONVERTOR OR TRANSLATOR NOT SPECIFIED

Code: S--Severe error
Explanation: Either the convertor or translator must be specified on the FILE statement if an input or output file is on 7-track tape.
System Action: The job ends after the generation phase.
User Response: Specify either convertor or translator on the FILE statement of the input or output file that is on 7-track tape.

SD176--INVALID ALTERNATE SEQUENCE STATEMENT

Code: S--Severe error
Explanation: You specified an alternate sequence in column 26 of your header line. The statement just read does not have ALTSEQ in columns 1-8. Also it is not a comment statement or a valid sort specifications statement.
System Action: The statement is ignored, and the job ends after the generation phase.
User Response: If the statement does not belong in the job, delete it. If it does belong in the job, correct it.

SD177--MISSING THE ** STATEMENT

Code: W--Warning
Explanation: Disk sort has read the ALTSEQ statements specified in column 18 of your header line. An ** (columns 1-2) statement should follow the last ALTSEQ statement that precedes your include/omit/field sort specification statements (I/O/F in column 6). However, the statement just read was an I/O/F statement.
System Action: Disk sort assumes an ** statement before the I/O/F statement.
User Response: Place an ** before the I/O/F statement before rerunning the job.

SD178--ALTSEQ STATEMENT HAS INVALID DATA

Code: S--Severe error
Explanation: The data to alter the normal sequence in the above ALTSEQ statement is in error. One of the 4-position groups (columns 9-12, 13-16, etc.) does not contain a valid hex number (0-9 or A-F).
System Action: The column(s) in error have been flagged with an * in the line printed above this message. Processing continues with the next 4-position group. The job ends after the generation phase.
User Response: Correct the statement and resubmit the job.

SD179--UNEXPECTED END-OF-FILE FOUND

Code: T--Terminal
Explanation: Disk sort was reading your ALTSEQ statements. Valid statements following an ALTSEQ statement are: another ALTSEQ statement, a comment statement, or an ** statement (columns 1-2). However, disk sort found an end-of-file statement. (If your sequence specification statements are being read from the system reader, the program read a /*, /&, or // END statement.)
System Action: Job ends at this point. No further statements are diagnosed.
User Response: Make sure the last part of your sequence specifications are not missing.

SD181--NO VALID ALTSEQ STATEMENTS FOUND

Code: W--Warning
Explanation: ALTSEQ is specified on the header line, but the system cannot find any valid ALTSEQ statements.
System Action: There will be no alternate collating sequence for this job; processing continues.
User Response: If you want an alternate collating sequence, cancel the job at halt C123 (Model 6) or halt 22 (Model 10, Model 12, or Model 15), correct the ALTSEQ statements, and resubmit the job. If you do not want an alternate collating sequence, remove the S from column 26 of your header line.

SD202—INVALID SPECIFICATION TYPE, COLUMN 6

Code: S—Severe error
 Explanation: The statement makes no sense to disk sort because it is not a comment statement, an include statement, an omit statement, or a field statement.
 System Action: Disk sort bypasses the statement. The job ends after the generation phase.
 User Response: This statement is either misplaced or incorrectly typed. Correct the statement and resubmit job.

SD204—BLANK CONTINUATION ASSUMED, COLUMN 7

Code: W—Warning
 Explanation: The referenced statement is the first in an omit or include set. However, it contains a nonblank character in column 7. A blank was expected.
 System Action: Disk sort ignores the nonblank character and processes the statement as though column 7 was blank.
 User Response: Correct statement before next job run.

SD206—INVALID CONTINUATION, COLUMN 7

Code: S—Severe error
 Explanation: The referenced statement is part of an include or omit set. It should have an A (AND) or O (OR) in column 7. However, it does not, nor does it contain a blank (see message SD208).
 System Action: Disk sort bypasses the statement and goes on to the next statement. Job will terminate at end of generation phase.
 User Response: Correct the statement before running the job again.

SD208—OR CONTINUATION ASSUMED, COLUMN 7

Code: W—Warning
 Explanation: The referenced statement is part of an include or omit set. It should contain an A (AND) or O (OR) in column 7. However, column 7 is blank.
 System Action: Disk sort assumes you want an O in column 7.
 User Response: Check the order of your statements. Make sure you want an O in column 7. If you do not, you can cancel the job at the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15).

SD210—INVALID C/Z/D/P/U/V SPECIFICATION, COLUMN 8

Code:	W—Warning	
Explanation:	Type of Statement	Column 8 Entry Should Be:
	Include or omit record type	C, Z, D, P, or U
	Control field statement ¹	C, Z, D, P, or U
	Data field or summary field	C, Z, D, U, P or V
System Action:	Disk sort assumes you want a C in column 8.	
User Response:	Check the character in column 8 and correct the statement before rerunning this job.	

¹ For an unconditional force or a force-all line, the column 8 entry must be a C.

SD212--SPECIFICATION IN WRONG LOGICAL ORDER

Code: S--Severe error
Explanation: If this is a control field statement (F in column 6), it improperly follows an omit statement (O in column 6). If this is an omit statement, it improperly follows an include statement (I in column 6).
System Action: This statement is bypassed and processing continues with next statement. The job ends after the generation phase.
User Response: Put your statement in proper order and resubmit the job.

SD214--INCLUDE OR OMIT STATEMENT AFTER INCLUDE-ALL

Code: S--Severe error
Explanation: Disk sort has previously read an include-all statement. Now only control field statements (F in column 6) should be read. However, an include or omit (I or O in column 6) was just read.
System Action: This statement and all following statements are processed as though the include-all card had not been read. The job ends after the generation phase.
User Response: Remove either the include-all statement or the include or omit statement. Resubmit job.

SD216--LAST SET NOT A VALID INCLUDE SET

Code: S--Severe error
Explanation: The last set of your sequence specifications should have been an include set. However, one of the following types of errors occurred:

1. There were no statements after the header statement.
2. The last set was an omit set.
3. The last set was an include set, but it contained no field statements.

System Action: The job will terminate at the end of the generation phase.
User Response: Make one of the following corrections depending upon the type of error that occurred:

1. Supply the sequence specifications.
2. Eliminate this omit set, because the records are omitted by default.
3. Supply the missing field statements. Resubmit the job.

SD218--TOO MANY SEQUENCE SPECIFICATIONS

Code: S--Severe error
Explanation: You have used all the main storage that was allocated by disk sort for the select/build routine and for error information.
System Action: Disk sort reads and error checks any following statements, but the job ends after the generation phase. The area allocated for the select/build routine and error information is now used only for error information.
User Response: Remove all causes of error messages. If this error occurs with no other error messages, decrease the number of sequence specifications.

SD220--TOO MANY ERRORS

Code: T--Terminal
Explanation: The main storage that disk sort allocated for the select/build routine and error information is now filled with error information (see message SD218).
System Action: Disk sort will not read or diagnose any more statements. The job ends at this point.
User Response: If your statements are being read from the system card input device, make sure you've included /* and /& statements.

SD222--ZONE OR V FIELD--LENGTH EXCEEDS 1

Code: W--Warning
Explanation: The referenced statement is a field statement (F in column 6). The type specified is a zone field (Z in column 8) or a V field (V in column 8) which should have a length of one byte.
System Action: The specified number of bytes (calculated using columns 9-12 and 13-16) is used in building the work record field for this record type. The byte of zone or overflow information generated is placed in the rightmost byte (low order byte) of this field. The other positions to the left contain binary zeros (control fields) or blanks (data fields).
User Response: If you do not want a length of one byte, respond to the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15) with a controlled cancel (option 2). The Halt Guide for each system lists all the halts and the options for each one.

SD224—COLUMNS 9-16 OR COLUMNS 20-27 ARE INVALID

Code: S—Severe error
Explanation: The location information in columns 9-16 and/or columns 20-27 of the referenced statement is invalid for one of these reasons:
1. The from location is greater than the to location.
2. The to location is blank or zero.
3. The lengths of factor 1 and factor 2 are not the same.
System Action: The job ends after the generation phase.
User Response: Correct this statement and resubmit the job.

SD226—DIGIT FIELD LENGTH EXCEEDS 16

Code: W—Warning
Explanation: A digit field longer than 16 bytes has been specified in the referenced statement.
System Action: Disk sort assumes you want a length of 16. The job ends after the generation phase.
User Response: If your digit field is longer than 16 bytes, divide it into two or more fields. Create new specifications for these fields and resubmit job. If your digit field is really 16 bytes or shorter, correct the statement and resubmit the job.

SD228—CHARACTER CONSTANT—LENGTH EXCEEDS 20

Code: W—Warning
Explanation: The referenced statement is an include or omit statement (I or O in column 6). Factor 2 contains a character constant since column 19 contains a C. The length implied in factor 1 (columns 9-12 and 13-16) is greater than 20, which is the maximum length allowed.
System Action: The rightmost 20 columns in factor 1 will be compared against the constant in columns 20-39.
User Response: If you do not want a 20-byte field you can end the job after the generation phase at the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15).

SD230—SAME SPECIFICATION TYPE ASSUMED, COLUMN 6

Code: W—Warning
Explanation: The referenced statement has a blank in column 6, so the specification type is not known. However, column 7 of this statement contains an A or O. The previous statement was either an include or omit statement (I or O in column 6).
System Action: Disk sort assumes the statement is a record type one with a specification type identical to that of the previous statement; that is, include (I) or omit (O).
User Response: If the statement is not a record type statement, you can cancel the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). In either case, correct the statement before re-running the job.

SD232—ZONE SPECIFIED—FACTOR 2 IS NOT A CONSTANT

Code: S—Severe error
Explanation: A zone was specified in column 8 of a record type statement (include or omit) and column 19 does not have a C (constant) specified.
System Action: Disk sort bypasses the statement. The job ends after the generation phase.
User Response: Correct the statement and resubmit the job.

SD234--ZONE SPECIFIED--RELATIONSHIP NOT EQ OR NE

Code: W--Warning
Explanation: A zone has been specified in column 8 of the referenced record type statement (include or omit) and columns 17-18 do not contain EQ or NE.
System Action: Disk sort assumes you want EQ in columns 17-18; processing of this statement continues.
User Response: If you do want an EQ in columns 17-18, you can continue the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). Correct the column 17-18 entry in the statement before the next job is run.

SD236--INVALID RELATIONSHIP, COLUMNS 17-18

Code: S--Severe error
Explanation: The referenced statement is a record type statement (include or omit) and columns 17-18 (the relationship between factor 1 and factor 2) does not contain an EQ, NE, LT, GT, LE, or GE.
System Action: System assumes an EQ entry in columns 17-18. The job ends after the generation phase.
User Response: Correct statement and resubmit job.

SD238--INVALID FACTOR 2 TYPE, COLUMN 19

Code: W--Warning
Explanation: The referenced statement is a record type specification (include or omit) and column 19 contains neither an F (field) nor C (constant).
System Action: Column 19 is assumed to contain a C; factor 2 is taken as a constant.
User Response: Check column 19 for an invalid character or missing character. If you do not want a C in column 19 you can cancel the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). In any case, you should correct the statement before running the job.

SD240--UNPACKED DECIMAL--LENGTH EXCEEDS 16

Code: W--Warning
Explanation: The referenced statement has a U in column 8. Columns 8-12 and 13-16 specify a factor 1 field longer than the allowable maximum of 16 bytes.
System Action: Disk sort processes the statement with a length of 16. The job ends after the generation phase.
User Response: Check the to and from columns and correct the statement before resubmitting job.

SD242--ZONE SPECIFIED--LENGTH EXCEEDS 1

Code: W--Warning
Explanation: The referenced statement is a record type specification (include or omit) with a Z (zone) in column 8. Factor 1 must be only one character long and it was not.
System Action: A length of 1 is assumed. Column 20 in this statement contains the constant that will be processed.
User Response: Verify that the assumption is correct. If it is not, you can cancel the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). In any case, you should correct the statement before rerunning the job.

SD244--PACKED DECIMAL--LENGTH EXCEEDS 8

Code: W--Warning
Explanation: Field length excessive. The P (packed) in column 8 dictates that the field must be only eight columns long (15 digits plus sign).
System Action: The specified number of bytes (calculated from columns 9-12 and 13-16) are used in building the work record field for this record type. Eight bytes of input information are used; the remainder of the field is unpredictable. The rightmost eight bytes of the field are used for include or omit. If column 19 contains a C, the constant in columns 20-35 of this statement is used for factor 2.
User Response: If you do not want to accept disk sort's assumptions, you can cancel the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). In any case, you should correct the statement before running the job again.

SD246--INVALID FIELD TYPE SPECIFICATION, COLUMN 7

Code: W--Warning
Explanation: The referenced statement is a field specification; however, column 7 does not contain an F, N, O, D or S.
System Action: The referenced field statement will be assumed as an N (normal control field) if there have been no previous data or summary data field specifications for this include set. Otherwise it will be assumed to be a D (data field) specification.
User Response: If you do not want an N in column 7, you can cancel the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). In any case, you should correct the statement before running the job again.

SD248--ILLEGAL FORCE-ALL CONTINUATION

Code: S--Severe error
Explanation: The referenced statement is a field statement for a force-all specification (F in column 7). However, it does not follow a conditional force specification as it should.
System Action: Disk sort ignores the statement. The job ends after the generation phase.
User Response: Correct statement and resubmit job.

SD250--FIELD STATEMENTS IN WRONG LOGICAL ORDER

Code: S--Severe error
Explanation: The referenced statement is a control field specification. However, a data (D) or summary (S) specification was previously specified in this include set. All control field statements (N, O, or F in column 7) must precede any data or summary data field specifications.
System Action: The statement is bypassed. Processing continues with next statement. The job ends after the generation phase.
User Response: This statement is either misplaced or incorrectly specified. Correct the statement and resubmit the job.

SD252--CONTROL FIELD LENGTH EXCEEDS HEADER VALUE

Code: S--Severe error
Explanation: The accumulated length of control field statements within the current include set exceeds length specified in header statement.
System Action: Disk sort processes all the following statements in this include set but the job ends after the generation phase.
User Response: Check the length specified in header line against the specified control fields. Correct the errors and resubmit job.

SD254—CONTROL FIELD LENGTH LESS THAN HEADER VALUE

Code: W--Warning
Explanation: The accumulated length of control fields specified for this include set is less than length specified in header statement.
System Action: All control field statements have been processed for this include set. The remaining positions are filled with binary zeros.
User Response: You can continue the job after the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15). Check control field length in header line before running the job again.

SD256—CHARACTER FIELD SPECIFICATION—LENGTH EXCEEDS 256

Code: S—Severe error
Explanation: The referenced statement is a character field specification (C in column 8). A length of more than 256 has been specified on this statement.
System Action: The job continues through the generation phases to diagnose the sort sequence specifications but terminates at the end of the generation phases.
User Response: If you require more than 256 positions of data, split the field into more than one field specification statement.

SD258—TO OR FROM FIELD NOT WITHIN INPUT RECORD

Code: S—Severe error
Explanation: Either or both specifications in the to or from columns are zero or contain a value larger than your input record length.
System Action: This job ends after the generation phase.
User Response: Check to make sure you have the desired input file. Also check that the from and to columns contain the desired values.

SD262—DATA LENGTH EXCEEDS HEADER VALUE

Code: S—Severe error
Explanation: The data field specifications (columns 6 and 7 contain FD or FS) for this include set now exceed your data length value. A SORTR or SORTRS job was specified. If you are dropping the control field (column 28 of the header line contains an X), the data length is the output record length (columns 29-32 of the header line). If you are not dropping the control field, the control field length (columns 13-17 of the header line) must be subtracted from the output record length to get the data field length.
System Action: The job ends after the generation phase.
User Response: There is an error in either the output record length in header line or the to and from fields in the field description lines. Correct the error and resubmit the job.

SD264—SPECIFICATIONS IN WRONG NUMERICAL ORDER

Code: W—Warning
Explanation: Columns 1-5 of each sequence specification statement are used to order the statements. One or more of the statements has a value in columns 1-5 that is less than the preceding statement. Disk sort prints an S to the left of the statements that are out of order.
System Action: Disk sort assumes the specification statements are in the order you want.
User Response: Check the order of the sequence specifications. If they are out of order, respond to the C123 halt (Model 6) or the 22 halt (Model 10, Model 12, or Model 15) with option 2. (The Halt Guide for each system explains the options following each halt.)

SD266--FACTOR 1 LENGTH EXCEEDS 256

Code: S--Severe error
Explanation: The referenced statement is an include or omit specification for character data (column 8 contains a C). Factor 1 is a field (column 19 contains an F). The length of factor 1 exceeds 256 bytes.
System Action: Disk sort assumes a factor 1 length of 256 but ends the job after the generation phase.
User Response: Break this statement into two specifications with an AND continuation line (put an A in column 7).

SD268--P OR U FACTOR 1 USED WITH ALTSEQ

Code: W--Warning
Explanation: The referenced include or omit record type statement specifies a packed or unpacked factor 1 and an alternate collating sequence is specified in the header statement.
System Action: Factor 1 and factor 2 are changed as indicated by the ALTSEQ statements. This change may affect the units position (and sign) of an unpacked decimal number or any one position of a packed decimal number. If it does, you may not include or omit the desired records.
User Response: Do not use P or U column 8 record type entries when you specify an alternate collating sequence.

SD269--PACKED FIELD USED WITH 96-COLUMN INPUT (Model 15 only)

Code: W--Warning
Explanation: 96-column cards can contain only the System/3 64-character set. Packed data cannot be contained on a 96-column card.
System Action: A 22 halt is issued at the end of the generation phase.
User Response: Verify that you are not trying to retrieve packed data from a 96-column card. The use of include sets may be necessary when using 96-column card input.

SD270--INVALID SPECIFICATION WITH TRANSLATE ON

Code: S--Severe error
Explanation: If input is on 7-track tape with TRANSLATE-ON, include, omit, or field specifications with packed data are not allowed. If output is on 7-track tape with TRANSLATE-ON, field specifications with packed data are not allowed.
System Action: The job ends after the generation phase.
User Response: Either correct the invalid specifications, change the output file to a 9-track tape file, or change the input file to a 7-track tape file with CONVERT ON specified.

SD271--INVALID SPECIFICATION UNLESS CONTROL FIELD DROPPED

Code: S--Severe error
Explanation: The control field is a normal packed or unpacked field, or an opposite packed, unpacked, zone, or character field. All of these control fields are modified by the disk sort program and will contain meaningless data.
System Action: The program ends after the generation phase.
User Response: Indicate that the control field is to be dropped (code an X in column 28 of the header statement) and repeat (as a data field) any of the control fields you want in the output record.

SD276--MORE THAN ONE SUMMARY V FIELD IN INCLUDE SET

Code: W--Warning
Explanation: This is a SORTRS job. The referenced statement specifies a summary V field (FSV in columns 6-8). At least one other summary V field has already been specified for this include set.
System Action: A data V field is assumed (FDV in columns 6-8).
User Response: Change the extra summary V field(s) to data V fields before running the job again.

SD278—SUMMARY V OR DATA FIELD INCONSISTENT

Code: W—Warning

Explanation: This is a SORTRS job. The referenced statement is a summary data field specification (S in column 7). Compare this specification with the first include set having S specifications; one of the following is true:

- The number of S specifications prior to this specification differs.
- The length of this S specification differs.
- The substitution character of this V specification differs.
- The relative location of the specification within the output record differs.

System Action: The S specifications of the first include set having S specifications are used to actually define the summary field positions.

User Response: A C123 halt (Model 6) or a 22 halt (Model 10, Model 12, or Model 15) will occur at the end of the Generation Phase. Check carefully to assure that the first include set with S specifications is correct. If not, correct and take a 2-option (controlled cancel).

SD280—NUMBER OF SUMMARY FIELDS INCONSISTENT

Code: W—Warning

Explanation: This is a SORTRS job. The total number of S specifications in this include set differs from the first include set having an S specification.

System Action: The S specifications of the first include set having S specifications are used to actually form the summary output records.

User Response: A C123 halt (Model 6) or a 22 halt (Model 10, Model 12, or Model 15) will occur at the end of the generation phase. Check carefully to assure that the first include set with S specification is correct. If not, correct and take a 2-option (controlled cancel).

SD282—NO SUMMARY SPECIFICATIONS FOUND—SORTRS JOB

Code: W—Warning

Explanation: A SORTRS job (summary tag along sort) was specified on the header statement for this job. However, no summary data field specifications were found in the statements for this job.

System Action: All records with duplicate control fields will be eliminated from the output file.

User Response: Check the job type to be sure that SORTRS was intended.

SD284—SUMMARY SPECIFICATIONS FOUND—SORTR JOB

Code: W—Warning

Explanation: This is a SORTR job. A summary data field specification (columns 6 and 7 contain FS) is expected only for SORTRS jobs. At least one S specification was found for this job.

System Action: All S specifications are treated as data specifications (column 7 contains a D rather than an S).

User Response: Check to assure that a SORTRS job was not wanted.

SD286—TOO MANY SUMMARY DATA FIELDS

Code: S—Severe

Explanation: This is a SORTRS job. More than 24 summary data fields (FS in columns 6 and 7 and C, D, P, or U in column 8) have been specified.

System Action: The job will end after the generation phase.

User Response: Limit the number of summary data fields to 24 or less.

SD288—INVALID SORTRS SPECIFICATION—
COLUMNS 20, 21, 22

Code: W—Warning
Explanation: The overflow field length in columns 20-22 is invalid, not right-justified, or less than the field length given by to and from entries for this specification.
System Action: Disk sort assumes that no entry was made in columns 20-22.
User Response: If a valid entry was meant for columns 20-22, cancel the job (C123 Halt—Model 6 or 22 halt—Model 10, Model 12, or Model 15), correct entries, and rerun the job. If no entry was meant for columns 20-22, the job will run but warning message will continue until error is corrected.

SD301—TAPE VARIABLE LENGTH BLOCK SPECIFIED

Code: T—Terminal
Explanation: Variable length blocks are not allowed on input files. Your input FILE statement is incorrect.
System Action: The job ends.
User Response: Disk sort accepts only fixed length block files from tape. Specify fixed block files on the FILE statement.

SD302—RECL AND/OR BLKL NOT SPECIFIED

Code: T—Terminal
Explanation: RECL and BLKL must be specified for a tape file on the FILE statement.
System Action: The job ends.
User Response: Specify RECL and BLKL and resubmit the job.

SD303—HEADER AND TAPE OUTPUT LENGTH DISAGREE

Code: T—Terminal
Explanation: The output record length specified on the output FILE statement and the output record length specified on the sort header statement are not the same.
System Action: The job ends.
User Response: Both record lengths (output FILE statement and sort header statement) must agree.

SD391—MAIN STORAGE ALLOCATED TOO SMALL

Code: T—Terminal
Explanation: The amount of main storage assigned to disk sort is not large enough to handle the job. On the Model 15, this halt can occur if the system was not generated with the support needed for a particular device such as tape, 5445, or card support. For CCP/Disk Sort, this occurs when the work record length is less than three bytes. This halt can also occur when a CCP/Disk Sort module is missing from the pack on which the CCP/Disk Sort resides.

System Action: This sort job cannot run. The CD45 halt (Model 6) or 25 halt (Model 10, Model 12, or Model 15) cancels the job.

User Response: Make sure you have the right input file. Check the SORTR or SORTRS output length specified in your header line if this is a tag along sort. Also check that the total number of bytes reserved for the sequence specifications (message SD461) does not exceed the recommended amount (see Appendix D). In addition, make sure there are no entries in positions 35-39 of the header statement (reserved for system use).

For CCP/Disk Sort, if the work record length is less than three bytes, then increase the control field record length or the output record length.

Ensure that all the CCP/Disk Sort modules are on the pack.

SD395—NOT ENOUGH SPACE FOR WORK FILE

Code: T—Terminal

Explanation: Disk sort automatic work file allocation is being used. One or more of the following situations has occurred:

1. The disk packs that are currently online cannot be used for work space because they contain multi-volume offline files or because a deferred mount is requested.
2. No work space is available on the online disk packs.
3. No work space is available on the disk packs indicated on the SWITCH statement (Model 12 and Model 15).

System Action: The job ends.

User Response: To make more scratch space available, perform one of the following:

1. Delete unnecessary files on the online disk packs (preferably the fixed units).
2. Mount a removable pack that has enough scratch space available.
3. Remove or modify the SWITCH statement (Model 12 and Model 15).

SD398—WORK FILE EXTENTS NOT USEABLE

Code: T—Terminal

Explanation: None of the track extents specified on the work FILE statement are large enough for a work block.

System Action: The job ends.

User Response: Increase the number of tracks on the work FILE statement so that the extents will be large enough to hold a work block.

SD401—JOB COMPLETED GENERATION PHASE

Code: I—Information
Explanation: This job has completed the generation phase. All diagnostics have been checked for errors.
System Action: If severe errors were found, message SD425 and halt CD45 (Model 6) or halt 25 (Model 10, Model 12, or Model 15) follow. If warning errors were found message SD422 and halt CD123 (Model 6) or halt 22 (Model 10, Model 12, or Model 15) follow. If no errors were found, the job proceeds to the execution phase of disk sort.
User Response: Review subsequent messages.

SD402—NO ERRORS FOUND

Code: I—Information
Explanation: No errors were found in the generation phase of disk sort.
System Action: Processing continues to the execution phase.
User Response: None.

SD403—OPERATOR CHOSE TO CANCEL JOB

Code: I—Information
Explanation: Message SD422 with halt CD123 (Model 6) or halt 22 (Model 10, Model 12, or Model 15) was just given. The warning error was not expected, so the operator chose to cancel the job.
System Action: Processing ends.
User Response: Correct the errors and resubmit the job.

SD404—OPERATOR CHOSE TO CONTINUE JOB

Code: I—Information
Explanation: Message SD422 with halt CD123 (Model 6) or halt 22 (Model 10, Model 12, or Model 15) was just given. The warning errors were expected, so the operator chose to continue the job.
System Action: Processing continues to the execution phase.
User Response: None.

SD422—XXX WARNING ERRORS HAVE BEEN FOUND 0—OK, CONTINUE 2—NOT EXPECTED, CANCEL

Code: A—Action
Explanation: No severe or terminal errors were found in the data supplied in your sequence specification statements. However, warning errors were found.
System Action: Halt C123 (Model 6) or halt 22 (Model 10, Model 12, or Model 15) occurs.
User Response: The operator should review the warning errors on the listing to decide what action to take, and then respond to the halt with option 0 (OK—continue the job) or option 2 (errors not expected—cancel the job).

SD423—PERMANENT I/O ERROR READING SPECIFICATIONS

Code: A—Action
Explanation: A permanent I/O error has occurred during the reading of your sort sequence specifications statements.
System Action: This disk sort job cannot continue. Halt CD3 (Model 6) or halt 23 (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3—Immediately cancel this job. Resubmit job. If error persists, contact IBM for hardware support.

SD425—SEVERE AND/OR TERMINAL ERRORS FOUND

Code: A—Action
Explanation: Severe and/or terminal errors were found during the generation phase of disk sort.
System Action: This job cannot continue. Halt CD45 (Model 6) or halt 25 (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3—Immediately cancel this job.

SD450--XX,XXX MAIN STORAGE BYTES ASSIGNED

Code: I--Information
Explanation: This is the size of main storage in which disk sort has been assigned to run.
System Action: None.
User Response: None.

SD451--XXXX BYTES--INPUT FILE RECORD LENGTH

Code: I--Information
Explanation: This is the input record length.
System Action: None.
User Response: None.

SD452--XXXX BYTES--WORK FILE RECORD LENGTH

Code: I--Information
Explanation: This is the record length for your work file. It was computed from the information on your header line.
System Action: None.
User Response: None.

SD453--XXXX BYTES--OUTPUT FILE RECORD LENGTH

Code: I--Information
Explanation: This is the record length for your output file. This information is taken from your header line.
System Action: None.
User Response: None.

SD454--XXXX BYTES--INPUT FILE BLOCK LENGTH

Code: I--Information
Explanation: This is the block length for your input file. This information is taken from your input FILE statement.
System Action: None.
User Response: None.

SD455--XXXX BYTES--OUTPUT FILE BLOCK LENGTH

Code: I--Information
Explanation: This is the block length for your output file. This information is taken from your output FILE statement.
System Action: None.
User Response: None.

SD461--XXXX BYTES--SELECT/BUILD ROUTINE

Code: I--Information
Explanation: The sequence specifications for this job (including alternate collating statements, if specified) will use the specified number of bytes of main storage.
System Action: None.
User Response: None.

SD462--XXX SEQUENCE SPECIFICATION STATEMENTS PROCESSED

Code: I--Information
Explanation: This is the number of sequence specifications statements (processed by the generation phase) in your job.
System Action: None.
User Response: None.

SD500—5445 FEATURE NOT AVAILABLE (Model 10 only)

Code: T—Terminal
Explanation: An attempt has been made to use a 5445 disk; however, the IBM System/3 Disk Sort 5445 Disk Storage Drive Feature is not on the program pack.
System Action: The job ends.
User Response: Either specify another device type or perform system generation with Disk Sort 5445 support.

SD501—DISK SORT TAPE SUPPORT NOT AVAILABLE (Model 10 only)

Code: T—Terminal
Explanation: An attempt has been made to use tape input or output; however, IBM System/3 disk sort tape support is not on the program pack.
System Action: The job ends.
User Response: Either specify another device type or perform system generation with disk sort tape support.

SD600—XX—FIRST PASS—#0 PASS—STARTED

Code: I—Information
Explanation: The execution phase of your job has begun. Records are being copied from the input file to the work file (XX = execution phase).
System Action: None.
User Response: None.

SD601—#XX PASS COMPLETED NNN

Code: I—Information
Explanation: This is simply status information. It tells you that pass #XX has just been successfully completed. (The first pass #00, the second pass is #01, etc.)
NNN is design information for the Model 15 (5704-SM9). For further explanation, see the diagnostic aids in the *IBM System/3 Disk Sort Program Logic Manual*, LY21-0517.
System Action: None.
User Response: None.

SD602—XXXXXX INPUT RECORDS WERE READ IN

Code: I—Information
Explanation: Number of records read from the input file.
System Action: None.
User Response: None.

SD603—NNNNNN RECORDS WERE SELECTED TO BE SORTED

Code: I—Information
Explanation: How many of the records just read in are going to be sorted.
System Action: None.
User Response: None.

SD605—CCP DATA TRUNCATED

Code: I—Information
Explanation: More than 80 bytes of data were passed or entered for the chain request. Only the leftmost 80 bytes of data were retained, the remainder was truncated.
System Action: None.
User Response: Review your task chain request to determine the impact of the truncation.

SD607—RELEASE/TASK CHAIN UNSUCCESSFUL

Code: I—Information
Explanation: The Release and Task Chain request could not be completed. Column 39 of the header statement contains a T indicating Release and Task Chain, but the sort program received control from the console or from a Task Chain.
System Action: A Task Chain operation will be attempted.
User Response: Check column 39 of the header statement and review your Release and Task Chain request. If symbolic files are used, the symbolic file reference will be lost.

SD62C—NO INPUT RECORDS WERE SELECTED TO BE SORTED

Code: A—Action
Explanation: Of the number of input records specified in message SD602, none of them met the INCLUDE requirements of your sort sequence specifications.
System Action: Halt C1234 (Model 6) or halt 2C (Model 10, Model 12, or Model 15) occurs.
User Response: Take option 0 to continue the job and create a null output file or take option 3 (immediate cancel). Review your sequence specifications in light of the records actually in your input file.

SD681—WORK RECORD COUNT IN ERROR

Code: T—Terminal
Explanation: Disk sort internal error.
System Action: The job terminates.
User Response: Collect pertinent information and contact IBM for programming support.

SD690—DISK SORT STARTING FINAL PASS

Code: I—Information
Explanation: The final pass of this sort job has begun. Records are being written on your output file.
System Action: None.
User Response: None.

SD900—TASK CHAIN UNSUCCESSFUL

Code: I—Information
Explanation: The task chain request could not be completed.
System Action: None.
User Response: Ensure that your task chain request is correct, and determine the impact of an unsuccessful chain on your job stream.

SD691—FINAL PASS SUCCESSFULLY COMPLETED

Code: I—Information
Explanation: The last pass of this sort was just successfully completed.
System Action: None.
User Response: None.

SD692—XXXXXX SORTED RECORDS PLACED ON OUTPUT FILE

Code: I—Information
Explanation: The last pass of this sort was just successfully completed. The output record selected and generated from the input records have been sorted and placed on your output file.
System Action: None.
User Response: None.

SD693—NO INPUT RECORDS FOUND

Code: A—Action
Explanation: No records were found in the input file.
System Action: Halt C1234 (Model 6) or halt 2C (Model 10, Model 12, or Model 15) occurs.
User Response: Check input file to ensure that it is the proper file and contains records you want to sort. Take option 0 to continue the job and create a null output file or take option 3 (immediate cancel).

SD694—XX PASSES REMAINING

NNN

Code: I—Information
Explanation: This is the number of passes remaining until job completion. NNN is design information for the Model 12 and Model 15. For further explanation, refer to the diagnostic aids in the *IBM System/3 Disk Sort Program Logic Manual*, LY21-0517.
System Action: None.
User Response: None.

SD695—XXXX USED FOR WORK FILE

Code: I—Information
Explanation: This is the number of tracks of work file used in this sort.
System Action: None.
User Response: None.

SD711—MOUNT OUTPUT TAPE

Code: A—Action
Explanation: The output tape is not required until the intermediate merge passes are done.
System Action: The program halts.
User Response: Mount the tape. If the output file uses the same tape drive as the input file, remove the input file and replace it with the output file reel. Do not reply to the halt until the output file reel is mounted.

SD901—*****NORMAL EOJ*****

Code: I—Information
Explanation: The sort job has been completed.
System Action: None.
User Response: None.

SD902—*****ABNORMAL EOJ*****

Code: I—Information
Explanation: Errors have prevented the sort job from being run.
System Action: None.
User Response: Review previous messages, correct errors, and resubmit job.

SD928—INPUT FILE PERMANENT I/O ERROR

Code: A—Action
Explanation: A permanent, unrecoverable I/O error has occurred during the reading in of records from your input file.
System Action: The job cannot continue. Halt C125 (Model 6) or halt 28 (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3 (immediate cancel). Resubmit the job. If error persists, contact IBM for hardware support.

SD929—WORK FILE PERMANENT I/O ERROR

Code: A—Action
Explanation: A permanent unrecoverable I/O error has occurred during the reading of records from your work file.
System Action: The job cannot continue. Halt CD4 (Model 6) or halt 29 (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3 (immediate cancel). Resubmit the job. If error persists, contact IBM for hardware support.

SD92A—OUTPUT FILE PERMANENT I/O ERROR

Code: A—Action
Explanation: A permanent unrecoverable I/O error has occurred writing to your output file.
System Action: The job cannot continue. Halt CD34 (Model 6) or halt 2A (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3 (immediate cancel) or option 2 (controlled cancel). Resubmit the job. If error persists, contact IBM for hardware support.

SD92D—SORT TASK CHAIN ATTEMPT FAILED

Code: A—Action
Explanation: After executing a CCP/Disk Sort program, the attempt to chain to the next task was not successful. Either not enough task chain queues were specified in the assignment set, or there is insufficient TP buffer space.
System Action: Message 2D (Model 15, 5704-SM7 only)
User Response: Option 1 (retry). Attempt to execute the task chain again. Option 2 (controlled cancel). The next program must be initiated by the operator.

SD92E--WORK FILE TOO SMALL

Code: A--Action
Explanation: While building the work file, disk sort ran out of space.
System Action: The job cannot continue. Halt C1245 (Model 6) or halt 2E (Model 10, Model 12, or Model 15) occurs.
User Response: Option 3 (immediate cancel). Allocate more tracks for the work file, and resubmit the job. If automatic work file allocation is being used, you can make more work area available by deleting unnecessary files (preferably on the fixed units), or by mounting a pack having adequate work area available. If automatic work file allocation is not being used, you can make more work area available by using a file statement designating adequate work area.

SD92F--OUTPUT FILE TOO SMALL

Code: A--Action
Explanation: When calculating the number of tracks needed for the output file, disk sort found that the output file was too small, or while writing to the output file, disk sort ran out of space.
System Action: Halt CD2 (Model 6) or halt 2F (Model 10, Model 12, or Model 15) occurs.
User Response: Option 1 (controlled cancel). If the halt occurs prior to the intermediate passes of the sort, another halt will be likely to occur in the later phases. A 1-option will allow the sort to take place and records to be put on the output file until no more space is available. Option 3 (immediate cancel). Either allocate more tracks for a disk output file or use a larger tape reel or multireel for a tape output file. Resubmit the job.

Appendix G. Instructions for Executing the Disk Sort Sample Program

When you receive the disk sort program, perform the following installation procedures:

1. If you received the disk sort program with the system control programs and have performed the program product generation during system generation, go to step 3.
2. For your system, perform the operations for Program Product Generation and Completing System Generation described in the *Operator's Guide*.
3. You should now verify that the disk sort program is operational by executing the disk sort sample program. General operating procedures for disk sort may be found in the Operator's Guide. To execute the sample program, perform the procedure under *Sample Program*; to execute it on the Model 6, perform the procedure under *Sample Program for Model 6*.

Note: There is no sample program for CCP/Disk Sort.

Sample Program for Model 6

For the Model 6, the disk sort sample program is a procedure on the distribution disk cartridge.

Mount distribution disk cartridge on R1 and set the DISK DRIVE 1 ON/OFF SWITCH to ON. When the disk drive is ready for use, move PROGRAM LOAD switch to ON.

You must now respond to the keywords that appear on the data entry keyboard.

Keyword	Your Keyboard Response
DATE—	mmddy (current month, date, and year) and press PROG START.
READER—	Press PROG START.
READY—	LOAD and press PROG START.
LOAD	NAME— \$D\$SPL and press PROG START.
	UNIT— R1 and press PROG START.
DATE	(mm/dd/yy)— Press PROG START.
SWITCH	(00000000)— Press PROG START.
FILE	NAME— INPUT and press PROG START.
	UNIT— R1 and press PROG START.
	PACK— Enter the proper pack identification and press PROG START.
	LABEL— Press PROG START.
	RECORDS— Press PROG START.
	TRACKS— 1 and press PROG START.
	LOCATION— Press PROG START.
	RETAIN— T and press PROG START.
	DATE— Press PROG START.

Keyword	Your Keyboard Response	Sample Program for Model 8, Model 10, Model 12, and Model 15
FILE	NAME— Press PROG START.	<p>For the Model 8, Model 10, Model 12, and Model 15, the disk sort sample program is a procedure on the distribution disk cartridge.</p> <p>With the system on F1, mount distribution disk cartridge on R1 and set the DISK DRIVE 1 ON/OFF SWITCH to ON. The READY light indicates when the disk drive is ready for use. If the system input device is other than the MFCU, set the ADDRESS/DATA switches as required. Set the PROGRAM LOAD SELECTOR to REMOVABLE DISK. Press PROGRAM LOAD, and enter the following statements via your system input device.</p> <pre> // DATE mmddyy // LOAD \$\$\$SPL, R1 // FILE NAME-INPUT, UNIT-R1, PACK-XXXXXX, // TRACKS-1 // RUN \$\$\$ RECL-0096, REC #000064 /& // CALL \$\$\$SPL, R1 // RUN </pre> <p>where mmddyy = correct date. XXXXXX = proper pack identification.</p> <p>These are all the entries that you need to make. As the sorted file is printed, the records should be in columns 10-15. The READY light indicates that the sample program has ended.</p> <p>For the Model 8 or Model 10, EJ is displayed when the file to be sorted has been built. To continue with the sort job, press console START or, if you have DPF, the program 1 HALT/RESET key.</p>
MODIFY	RUN and press PROG START. \$\$\$ RECL-0096, REC#-000064 and press PROG START.	
READY—	CALL and press PROG START.	
CALL	NAME— \$\$\$SPL and press PROG START. UNIT— R1 and press PROG START.	

These are all the responses that you must make. When the sorted file is printed, the records should be in descending order in columns 10-15. When the READY keyword appears, the sample program has ended.

```

// DATE mmddyy
// LOAD $$$SPL, R1
// FILE NAME-INPUT, UNIT-R1, PACK-XXXXXX,
// TRACKS-1
// RUN
$$$ RECL-0096, REC #000064
/&
// CALL $$$SPL, R1
// RUN

```

where
mmddyy = correct date.
XXXXXX = proper pack identification.

These are all the entries that you need to make. As the sorted file is printed, the records should be in columns 10-15. The READY light indicates that the sample program has ended.

For the Model 8 or Model 10, EJ is displayed when the file to be sorted has been built. To continue with the sort job, press console START or, if you have DPF, the program 1 HALT/RESET key.

Appendix H. CCP/Disk Sort Specifications Differences

This appendix is designed as a reference guide for the following System/3 CCP/Disk Sort users:

- Model 4, Program Number 5703-SM2
- Model 15, Program Number 5704-SM7

CCP/DISK SORT OVERVIEW

The IBM System/3 CCP/Disk Sort Program is a disk-resident program which sorts a data file into either ascending or descending sequence. The sorting is executed under the control of the Communications Control Program (CCP) on System/3 Models 4 or 15. Any disk file organization supported by the System/3 disk system management may be used as input to the sort. This includes indexed, direct, and sequential files.

The user supplies OCL statements and sort specifications to the CCP/Disk Sort generator program (\$DGSRT). \$DGSRT does not run under control of CCP; rather, it generates an object program that can later be executed as a user task under CCP control. This same object program can also execute in a batch mode, rather than under CCP. In order to execute the object program, either in batch mode or under CCP, the entire CCP/Disk Sort Program (\$DG.ALL) must be present on the program pack during generation and execution.

The CCP/Disk Sort Program requires the use of the Overlay Linkage Editor during the generation phase. The Overlay Linkage Editor must reside on the same pack on which the CCP/Disk Sort Program resides.

OCL statements for \$DGSRT can be supplied from the system input device or from a procedure library member. Sort specifications can be supplied from the system input device, from a procedure, or from a source library member.

A listing of the sort specifications and informational messages is optional. (See column 27 in Chapter 3.) When requested, they are listed on the printer only if log on is specified.

The number of sort specifications may be more limited when generating a CCP/Disk Sort program because it must run in a minimum of 12 K bytes of main storage.

FILE STATEMENT CONSIDERATIONS

For CCP/Disk Sort generation, the FILE statements for input, work, and output are needed to provide the LABEL parameters. The LABEL parameters are used as the file NAME for the object program execution. This is done so that multiple sorts can be executed in a single CCP session.

If any of the sort files (input, work, and output) are physically online during generation, the FILE statement must be valid for that file.

If the sort file is not online the FILE statement must still be syntactically correct and contain a TRACKS parameter and a valid LABEL parameter. The UNIT and TRACKS parameters do not have to be the actual values to be used when running the sort, but they must be present. The UNIT parameter cannot specify a nonremovable unit.

\$\$SOURCE and \$WORK files must be provided for the generation of the object module, and both files must reside on a 5444 simulation area.

Multiple Input Files (Model 15 only)

Input can be from either of the following devices:

- 3340 Direct Access Storage Facility
- 3344 Direct Access Storage

Up to eight files can be used as input using any combination of the preceding devices. The file names for generation must be INPUT or INPUT1 (but not both), INPUT2, INPUT3, INPUT4, INPUT5, INPUT6, INPUT7, and/or INPUT8¹.

The files are processed INPUT (or INPUT1) serially to INPUT8 regardless of the order of the OCL FILE statements. The lowest numbered input file is processed fully first, then the next numbered input file will be processed in its entirety, then the next until all the files have been individually processed in sequence. For a particular disk sort run, the input files do not have to be numbered serially; that is, the file names may be INPUT3, INPUT5, and INPUT8¹, or a particular run could have just INPUT8 as the input file.

Note: The record lengths of all disk files must be the same. Also note that the UNIT parameter is used.

¹ The LABEL parameters are used as the file NAME for the object program execution.

Output File Over Input File Considerations for Model 4

If Input File Is	Output Over Input?
Direct (See Note)	No
Indexed (See Note)	No
Sequential	Yes
<i>Note:</i> Addrout sort is recommended.	

Work Record Length

Work record lengths cannot be less than three bytes. This means that the length of the sort control field, plus the length of the data fields (if any), must be three bytes or greater. If the work record length is less than three bytes, a diagnostic message (SD391) will be issued during sort generation and the sort cannot be executed. If necessary, specify a minimum 3-byte control field which you can drop (X in column 28 of the header statement) at output time. Be careful, there may be data in portions of the control field that may be required in the output record.

Note: The work record length discussed is not the value specified in the CCP DISKFILE statement. The record length specified in the DISKFILE statement for the work file must always be 256.

CCP/Disk Sort Object Program Generation Job Stream
Example

```

      1  2
// LOAD $DGSRT,unit

      3
// COMPILE OBJECT-XX, (LINK ADD-8000)

      4
// FILE NAME-$SOURCE,TRACKS-10, . . .

      4
// FILE NAME-$WORK,TRACKS-10, . . .

      5          6
// FILE NAME-INPUT,LABEL-SAMPIN, . . . (file description)

      7          8          6
// FILE NAME-WORK,LABEL-SAMPWK, . . . (file description)

      9          6
// FILE NAME-OUTPUT,LABEL-SAMPOT, . . . (file description)

// RUN

    10
(header specification)

    11
(record type specifications)

    12
(field specifications)

/* (end of file) }
// END           } Use either
                  } statement
```

1 \$DGSRT is the CCP/Disk Sort module name.

2 The unit parameter specifies the pack on which the CCP/Disk Sort program resides. The codes are F1, R1, F2, or R2.

3 The COMPILE statement is valid for the Model 15 only. The OBJECT parameter specifies the unit on which the object program is to be catalogued.

4 The overlay linkage editor requires that the \$SOURCE and \$WORK files be 10 tracks.

5 The LABEL keyword identifies the actual input file name at execution time (infile — see item **1** of *CCP Assignment Set Statement Example*).

6 Same as for Disk Sort (see the discussion on *OCL Statements* in Chapter 2). Also see the heading *File Statement Considerations* in this appendix.

7 The FILE statement for WORK must be present in CCP/Disk Sort.

8 The LABEL keyword identifies the actual work file name (wrkfil — see item **2** of *Example, CCP Assignment Set Statements*).

9 The LABEL keyword identifies the actual output file name (outfil – see item **3** of *CCP Assignment Set Statement Example*).

10 The header specification field definitions for columns 1-34 for CCP/Disk Sort are the same as those defined for disk sort (see Chapter 3).

CCP/Disk Sort header specification field definition variations for columns 35-80 follows.

Columns 35-36 are reserved for system use. Column 35 cannot be used because only one sort algorithm exists.

Columns 37-38 contain the Model 15 storage size required to execute the object module. Specify a value from 12K to 48K in 2K increments.

Column 39: Model 4 must contain an R, which indicates that a record length exists in columns 40-43. Model 15 must contain either a C or an R. A C indicates that the sort chains to the next task. An R or C indicates that a record length exists in columns 40-43.

Note: For program numbers 5799-ATF, 5799-ATG, and 5799-ATH, column 39 can contain only an R.

Columns 40-43 contain the input file record length, right justified with or without leading zeros.

Note: If a user comment exists in columns 40-43, you must either shorten the comment or remove it.

Columns 44-72 contain any comment you wish to use.

Note: If you are redoing the header statement for a program not previously used by CCP/Disk Sort, remember to rewrite the comment since columns 40-43 now contain data of a different nature.

Columns 73-74 are not used.

Columns 75-80 contain the name given to the object program. This is the name that you enter from the terminal to invoke the sort. The entry can be:

1. A valid program name. The first character must be alphabetic, but cannot be a #, \$, or @ character. The remaining characters must be alphanumeric with no embedded blanks or special characters.
2. Blank. The default is SRTOBJ.

Note: DIR, ALL, and SYSTEM are reserved names and must not be used as program names.

11 For record type specifications, see Chapter 4.

12 For field specifications, see Chapter 5.

CCP Assignment Set Statements Example

```

      1
// DISKFILE NAME-SAMPIN,ORG-C,RECL-aaaa

      2
// DISKFILE NAME-SAMPWK,ORG-C,RECL-256

      3
// DISKFILE NAME-SAMPOT,ORG-C,RECL-bbbb

// PROGRAM NAME-anynam,PGMDATA-YES,SORT-YES,
      1          2          3
FILES-'SAMPIN/CG/NOSHR,SAMPWK/CA/NOSHR,SAMPOT/CO/NOSHR'
```

- 1** SAMPIN (infile) is the same LABEL keyword given for FILE NAME-INPUT in the job stream (see item **3** in the *Sample CCP/Disk Sort Job Stream*). This file can be SHR or NOSHR and the value of aaaa is the same value used in columns 40-43 of the header statement (record length).
- 2** SAMPWK (wrkfil) is the same LABEL keyword given for FILE NAME-WORK in the job stream (see item **6** in the *Sample CCP/Disk Sort Job Stream*). This file must be NOSHR and RECL must be 256.
- 3** SAMPOT (outfil) is the same LABEL keyword given for FILE NAME-OUTPUT in the job stream (see item **7** in the *Sample CCP/Disk Sort Job Stream*). This file must be NOSHR. The value of bbb can be:
 - The value specified in columns 29-32 of the header statement.
 - A 3 if SORTA is specified.

CCP/Disk Sort Execution (Batch Mode) Job Stream
Example

```

1
// LOAD name,unit

                               2
// FILE NAME-SAMPIN,UNIT-xx,PACK-xxxxxx,...
                               2
// FILE NAME-SAMPWK,UNIT-xx,PACK-xxxxxx,...
                               2
// FILE NAME-SAMPOUT,UNIT-xx,PACK-xxxxxx,...

// RUN
```

- 1** The name is the name given to the object program in columns 75-80 of the header card at generation time. If no name was specified, the default is SRTOBJ.
- 2** These parameters must all be correct for the input, work, and output files.

Appendix I. CCP/Disk Sort Program for System/3 Models 8, 10, 12, 15A, 15B, and 15C

Appendix I points out the differences between the CCP/Disk Sort Programs for Models 8, 10, 12, and 15A, 15B, and 15C and the corresponding Disk Sort Program for these systems. For a description of the CCP/Disk Sort Program, refer to Appendix H of this manual.

These CCP/Disk Sort Programs are classified as programming RPOs (request for price quotation) rather than as Program Products. Your IBM marketing representative can provide information about RPOs.

Note: Automatic work file allocation is not supported by CCP/Disk Sort Programs.

CCP/DISK SORT PROGRAM (PROGRAM NUMBER 5799-ATF)

This CCP/Disk Sort Program is used on System/3 Model 8 or 10, and operates under control of SCP 5702-SC1. The generated object program can be executed under control of the Communications Control Program (program number 5702-SC1, feature 6033). The CCP/Disk Sort Program (program number 5799-ATF) is functionally identical to the Disk Sort Program for the Models 8 and 10 (Program Product number 5702-SM1).

Using CCP/Disk Sort, sort generation requires 12K bytes of main storage and a 5444 Disk Storage Drive.

To execute the generated sort program, 12K bytes of main storage is required. The following devices are supported for the input, work, and output files:

- 5444 Disk Storage Drive
- 5445 Disk Storage
- 5448 Disk Storage Drive

CCP/DISK SORT PROGRAM (PROGRAM NUMBER 5799-ATG)

This CCP/Disk Sort Program is used on System/3 Model 12 and operates under control of SCP 5705-SC1. The generated object program can be executed under control of the Communications Control Program (program number 5705-SC1, feature 6070 or 6071). The CCP/Disk Sort Program is functionally identical to the Disk Sort Program for the Model 12 (Program Product number 5705-SM1) with one exception, only one input file can be specified on this CCP/Disk Sort Program.

Using CCP/Disk Sort, sort generation requires 12K bytes of main storage and a 3340 Direct Access Storage Facility.

To execute the generated sort program, 12K bytes of main storage is required. Input, work, and output files can reside in the main data area or in the simulation area of a 3340 Direct Access Storage Facility.

CCP/DISK SORT PROGRAM (PROGRAM NUMBER 5799-ATH)

This CCP/Disk Sort Program is used on System/3 Model 15A, 15B, or 15C and operates under control of SCP 5704-SC1. The generated object program can be executed under control of the Communications Control Program (program number 5704-SC1, feature 6033, 6070, or 6071). The CCP/Disk Sort Program is functionally identical to the Disk Sort Program for the Model 15 (Program Product number 5704-SM1) with one exception, the only devices supported by the CCP/Disk Sort Program for the input, work, and output files are the following:

- 5444 Disk Storage Drive
- 5445 Disk Storage
- 3340 Direct Access Storage Facility

In using CCP/Disk Sort, sort generation requires 12K bytes of main storage.

To execute the generated sort program, from 12K to 48K bytes can be specified; however, the maximum task size under CCP is 32K bytes.

ADDROUT sort: The type of sort where the output consists of record addresses in the input file.

alphabetic character: Letters A through Z and special symbols #, \$, and @.

alphanumeric character: The alphabetic characters defined and numeric characters 0-9.

block: A predefined unit of storage space for files on the disk. For example, 1 block = 2560 bytes.

byte: The representation of a character by eight binary digits; the amount of storage required for one EBCDIC character.

conditional force: The control field character is replaced before the record is sorted, only if the control field in the input record contains a particular entry.

constant: Data that does not change during execution of a program.

control field: One or more specified fields that are compared to determine the record sequence in the output file.

data: A collection of facts, numbers, letters, and symbols that can be processed by a computer.

data field: See *field*.

execution phase: The part of the sort program that gets the records you want to sort, sorts the records, and writes the sorted records onto the output file.

field: One or more adjacent record positions that contain related information.

file: An organized collection of related records.

file description: The specification for the arrangement of fields within a record and records within a file.

forced control field: A one-position control field that results from replacing the value or adding a value.

generation phase: The part of the sort program that translates the sequence specifications into machine language.

input: Data transferred to computer storage for processing.

integer: A whole number stored in hexadecimal (binary) form.

library: An area on disk that contains procedure members, source members, load members, and subroutine members as well as areas required by the system control program. For example, sort is stored in the library.

normal control field: Any control field that is specified by an N in column 7 of the field sequence specifications.

OCL: Operation control language.

OCL statement: A statement that identifies the job and its processing requirements to the control program.

opposite control field: Any control field that is specified by an O in column 7 of the field sequence specifications.

output: Data transferred from computer storage to printed form, disk, or to the display screen.

output file: File resulting from the sort.

overflow field: In a summary tagalong sort, the field that allows for anticipated field expansion.

packed data field: Each byte within a field represents two numeric digits except the low-order byte, which contains the sign. For non low-order bytes, bits 0-3 for one digit; bits 4-7 for one digit.

procedure: A named collection of related OCL statements, and possibly control statements that perform a particular task.

procedure member: A procedure stored in the library.

program date: The date associated with a program (job step).

real number: A number stored in IBM System/360 floating-point format, in which the leftmost byte contains the sign and exponent (decimal location) and the remaining bytes contain the fraction or mantissa in left-justified binary form.

record: A collection of related data, treated as a unit. For example, one line of an invoice might form a record.

relative record number: In a direct file, the location of a record in relation to the beginning of the file.

right-justify: The placement of data in a field with the last significant byte in the rightmost position.

sequence specification. Defines the arrangement of a file.

source member: A collection of records used as input to a program such as RPG II specifications or sort sequence specifications. Source members are stored in the library.

summary data field. In a summary tagalong sort, a data field designated for accumulated totals.

summary tagalong sort: A type of sort that arranges records by a specified control field and accumulates totals from specified fields within those records.

tagalong sort: A type of sort that arranges records by a specified control field and includes other specified data fields in the output.

unpacked data field: Each byte within a field represents one numeric digit. Bits 0 through 3 are reserved for the sign; bits 4 through 7 are the numeric portion.

work file: An area reserved on the disk for temporary storage of the data being sorted.

*
to indicate comments on field description line 37
to indicate comments on record type lines 30
// SOURCE statement
in BUILD cycle for Model 6 98
in CALL cycle for Model 6 99
in LOAD cycle for Model 6 94
to supply sequence specifications for Model 10 and
Model 15 103
//SWITCH statement, directed automatic work file
allocation 14

action message (A)
explanation of 121
printed by the disk sort program 24
address out sort (addrout), definition and function 5
addrout and tag along sorts, comparison of 14
addrout sort (SORTA) (see SORTA)
alphameric constants, as factor 2 (columns 20-39--record
type specifications) 35
alphameric data
as affected by column 8 record type entry 30
effect on column 8 record type entry 31
alternate collating sequence
ALTSEQ statement use 107
effect on factor 1 and factor 2 comparison 34
effect on sort time 111
indicated on header line (column 26) 23
restrictions on using column 8 record type entries 107
ALTSEQ statements
examples of 108
how to code 107
placement in sequence specifications 108
programming considerations 108
what they can and cannot change 108
AND lines
definition 29
describing a record type 28
describing additional factor 1 fields 34
in include sets 28
in omit sets 29
ASCII file 10, 11, 12
automatic work file allocation
directed 5445, Models 12 and 15 13, 14
effect on sort time 113
file considerations 13, 14
multiple file input 10, 11, 12

binary relative record number 5
BUILD cycle for Model 6
function 91
keyword-response summaries 95
when to use 91
bytes in a track, number of
3340, for the 116
5444, for the 116
5445, for the 116

C/Z/D/P/U (column 8--record type specifications)
effect on factor 1 length (columns 9-16) 34
restrictions when using alternate collating sequence 31
to indicate packed or unpacked data 30
to indicate portions of characters to compare 31
to indicate records you want to sort 31
when factor 1 and factor 2 contain alphameric data 31
with signed and unsigned data 31
C/Z/D/P/U/V (column 8--field description specifications)
effect on your sorted file 40
what it tells disk sort 40
calculating file size 116
CALL cycle for Model 6
function 91
keyword-response summaries 99
when to use 91
CCP/disk sort
\$SOURCE 150
\$WORK 150
field description specifications (see field description
specifications)
FILE statement considerations 150
header specifications
column summary 22
column 6 (line type) 23
column 18 (sequence) 23
column 26 (collating sequence) 23
column 27 (print option) 24
column 28 (output option for SORTR) 24
column 34 (nonverify option) 24
columns 37-38 (storage size) 25, 151
column 39 (record length code) 25, 151
columns that you must fill in 22, 151
columns 1-2 (page numbers) 23
columns 3-5 (line number) 23
columns 7-12 (job) 23
columns 13-17 (largest total of control fields of any
record type) 23
columns 29-32 (output record length for SORTR) 24
columns 40-43 (record length) 25, 151
columns 44-72 (job description) 25, 151
columns 75-80 (program name) 25, 151

CCP/disk sort (continued)

- overview 148
- record type specifications (see record type specifications)
- sample CCP assignment set statements 152
- sample job stream 150
- specifications, differences in 150
- work record length, specifying 150

characters

- in compare operations 30
- representation within the computer 31

collating sequence (column 26—header specifications)

- alternate 107
- specified for compare operations 23
- standard 107
- table of 106

column 5 (field description specifications), to add new sequence

- lines without renumbering 37

column 5 (record type specifications), to add new sequence

- lines without renumbering 26

column 6 (line type)

- field description specifications 37
- header specifications 22
- record type specifications 26

column 7 (continuation or comments), record type specifications 30

column 7 (field type or comments), field description specifications 38

column 8 (C/Z/D/P/U)

- record type specifications 30

column 8 (C/Z/D/P/U/V)

- field description specifications 41

column 17 (conditionally forced character), field description specifications 45

column 18

- field description specifications (forced character) 45
- header specifications (sequence) 22

column 19

- field descriptions (specifying replacements) 46
- record type specifications (field or constant) 34

column 26 (collating sequence), header specifications 22

column 27 (print option), header specifications 22

column 28 (output option), header specifications 22

column 34 (nonverify option), header specifications 22

column 39 (CCP/disk sort record length code), header specifications 25, 151

columns 1-2 (page number)

- field description specifications 37
- header specifications 22
- record type specifications 26

columns 3-5 (line number)

- field description specifications 37
- header specifications 22
- record type specifications 26

columns 7-8 (field description specifications), possible combinations of entries 37, 43

columns 7-12 (job), header specifications 22

columns 9-12 (from entry), field description specifications 37, 44

columns 9-16 (factor 1 location), record type specifications 26, 34

columns 13-16 (to entry), field description specifications 37, 44

columns 13-17 (largest total of control fields of any record type), header specifications 22

columns 17-18 (relation), record type specifications 26, 34

columns 20-22 (overflow field length entry), field description specifications 37, 52

columns 20-27 (factor 2 field), record type specifications 26, 35

columns 20-39 (factor 2 constant), record type specifications 26, 35

columns 29-32 (output record length), header specifications 22, 24

columns 40-43 (CCP/disk sort record length), header specifications 25, 138

columns 40-72 (comments)

- field description specifications 37, 52
- record type specifications 26, 36

columns 40-72 (job description—disk sort), header specifications 22, 25

columns 44-72 (CCP/disk sort job description), header specifications 25, 151

columns 75-80 (program identification/program level), header specifications

- CCP/disk sort 25, 151
- disk sort 25

comment line

- field description specification 38
- record type specifications 30

comments (columns 40-72—field description specifications) 37, 52

comments (columns 40-72—record type specifications), to describe record types 26, 36

compare operations

- affected by column 8 of record type specifications 30
- affected by column 19 of record type specifications 35
- affected by columns 17-18 of record type specifications 34

conditional force

- explanation of 39
- using a force-all line 51
- using normal or opposite and forced control fields 49
- using only forced control fields 48

conditionally forced character (column 17—field description specifications)

- examples 48
- how to define 46
- used to assign value to summary overflow field 46
- what it tells disk sort 45

constant

- as factor 2 (columns 20-39—record type) 35
- indicated by column 19 of record type specifications 35
- length when used as factor 2 on record type specifications 35
- position when used as factor 2 on record type specifications 35
- positioning sign in numeric constants 35

continuation or comment (column 7—record type specifications)

- to describe different record types having identical field types 30
- to further describe the same record type 30
- to indicate comment line 30
- to indicate start of include or omit set 30

control and data fields, when and how to describe one field as both 24

- control fields
 - forced 39
 - how order affects sort time 113
 - how to drop 24
 - normal 39
 - number of when two or more record types 38
 - opposite 39
 - packed or unpacked 41
 - specifying use by disk sort 38
 - to sequence information in sorted records 42
 - total lengths when two or more record types 38
 - types of 39
 - when to drop 41
- conversational OCL for Model 6
 - interrelationship with sequence specifications 101
 - summary of keywords and responses for OCL cycles 92
 - to supply OCL statements 9, 16
 - when to use each OCL cycle 91
- cycles, OCL
 - functions of 91
 - list of 91
- data
 - alphanumeric 31
 - numeric 32
 - representation within the computer 31
 - used in compare operations 30
- data and control fields, when and how to describe one field as both 24
- data fields
 - definition 38
 - describing on field description line 38
 - number of when two or more record types 38
 - total lengths when two or more record types 38
- data fields, summary (see summary data fields)
- data recorder, to supply OCL statements 16
- defining
 - conditional force characters 46
 - force-all characters 46
 - unconditional force characters 46
- diagnostic messages print option 24
- digit portion of System/3 character 31
- disk address 5
- disk sort program
 - execution phase 121
 - files used 6
 - general description 3
 - generation phase 121
 - information OCL statements give your computer 9
 - input needed 6
 - interaction with input, work, and output files 4
 - timings, sample 106
- dropping control fields
 - how to do it 22, 24
 - when to do it 41
- error messages
 - explanation of codes 121
 - list of 122
- example of using force 63
- examples of addroot and tag along sorts 55, 72
- examples of integer and real number sorting 83
- examples of summary sort jobs 67
- examples of using forced control fields 47
- execution phase 121
- external indicators 14
- factor 1 (columns 9-16—record type specifications)
 - changed by ALTSEQ statements 108
 - coding rules 34
 - coding when only one character 34
 - describing more than one factor 1 field 34
 - how length affected by column 8 entry 34
 - possible lengths 34
 - to identify included or omitted records 34
- factor 2 (columns 20-27 or 20-39—record type specifications)
 - changed by ALTSEQ statements 108
 - coding the sign in packed and unpacked numeric constants 35
 - how to position the numeric constant in columns 20-39 35
 - length of factor 2 constant 35
 - length of factor 2 field 35
 - rules for coding 35
- factor 2 field (columns 20-27—record type specifications)
 - coding rules 35
 - coding when one character long 35
 - relationship to factor 1 field 35
- field description specifications (disk sort and CCP/disk sort)
 - as part of include set 28
 - column summary 37
 - column 6 (line type) 38
 - column 7 (field type or comments) 38
 - column 8 (C/Z/D/P/U/V) 40
 - column 17 (conditionally forced character) 45
 - column 18 (forced character) 45
 - column 19 (specifying replacements) 46
 - columns you must fill in 37
 - columns 1-2 (page number) 38
 - columns 3-5 (line number) 38
 - columns 9-12 (from) 44
 - columns 9-16 (field location) 44
 - columns 13-16 (to) 44
 - columns 20-22 (overflow field length entry) 52
 - columns 40-72 (comments) 52
 - explanation of the column entries 38
 - what they describe for addroot sorts 38
 - what they describe for summary sort 38
 - what they describe for tag along sorts 38
 - what they tell disk sort 37
 - when using an include-all line 27
- field length, dependent on column 8 entry (field description specifications) 45
- field location (columns 9-16—field description specifications)
 - to identify location of field in input record 44

field or constant (column 19—record type specifications) to determine if factor 2 is constant or field 35

field type or comments (column 7—field description specifications) 38

file considerations 10

file location on disk, effect on sort time 111

FILE statement considerations 150

force-all

- explanation 39
- result if not used 46
- used after conditional force lines 40

force-all character

- examples of 51
- how to define 46

force, example of 63

forced character (column 18—field description specifications)

- used to initialize summary overflow field 46
- when and how character is forced 45

forced control fields

- examples 47
- rules for using 42
- types 39
- when to use 39

format of records

- definition 8
- effect on order of sequence specifications 8

from entry (columns 9-12—field description specifications)

- describing one character field 44
- to indicate first position of field in input record 44

from entry (columns 9-12—record type specifications) to indicate first position of factor 1 in input record 34

from entry (columns 20-23—record type specifications) to indicate first position of factor 2 field in input record 35

generation phase 110

header specifications, CCP/disk sort

- column summary 22
- column 6 (line type) 23
- column 18 (sequence) 23
- column 26 (collating sequence) 23
- column 27 (print option) 24
- column 28 (output option for SORTR) 24
- column 34 (nonverify option) 24
- column 39 (record length code) 25, 151
- columns that you must fill in 22, 151
- columns 1-2 (page numbers) 23
- columns 3-5 (line number) 23
- columns 7-12 (job) 23
- columns 13-17 (largest total of control fields of any record type) 23
- columns 29-32 (output record length for SORTR) 24
- columns 40-43 (record length) 25, 151
- columns 44-72 (job description) 25, 151
- columns 75-80 (program name) 25, 151

header specifications, disk sort

- column summary 22
- column 6 (line type) 23
- column 18 (sequence) 23
- column 26 (collating sequence) 23
- column 27 (print option) 24
- column 28 (output option for SORTR) 24
- column 34 (nonverify option) 24
- columns that you must fill in 22
- columns 1-2 (page numbers) 23
- columns 3-5 (line number) 23
- columns 7-12 (job) 23
- columns 13-17 (largest total of control fields of any record type) 23
- columns 29-32 (output record length for SORTR) 24
- columns 40-72 (job description) 25
- explanation of column entries 23
- what they tell disk sort 22

include line, to select records to sort 27

include set

- definition 28
- mixing with omit sets 28
- rules for using 28
- when to use 28

include-all line

- implied 8
- to sort records that have not been included or omitted 27

information message (I) 121

initializing an overflow field 52

input file

- calculating storage tracks for 116
- considerations 9
- considering overlay when programming 13
- defined during BUILD cycle 84
- defined during LOAD cycle 81
- how its position on disk affects sort time 112

input sequence specifications, order when using ALTSEQ 108

input to disk sort, file types that you can sort 6

instructions for executing disk sort sample program

- for Model 6 146
- for Model 10 147
- for Model 15 147

job (columns 7-12—header), to specify the type of sort 23

job description (columns 40-72—header specification for disk sort), to supply comments 25

job description (columns 44-72—header specification for CCP/disk sort), to supply comments 25, 151

- largest total of control fields of any record type (columns 13-17--header specification), calculating entry length 23
 - factor 1 of record type specifications 34
 - factor 2 constant (columns 20-39--record type specifications) 35
 - factor 2 field (columns 20-27--record type specifications) 35
 - maximum field length for columns 7-8 (field description specifications) 43
- library source 103
- line number (columns 3-5) with page number (columns 1-2) for sequencing
 - field description specifications 38
 - header specifications 23
 - record type specifications 27
- line type (column 6)
 - to identify field description line (field description specifications) 38
 - to identify header line (header specifications) 23
 - to identify include line (record type specifications) 27
 - to identify include-all line (record type specifications) 27
 - to identify omit line (record type specifications) 27
- LOAD cycle for Model 6
 - function 91
 - keyword-response summary 92
 - when to use 91
- main storage size, effect on sort time 110
- main storage, amount used by sequence specifications 111
- messages
 - alphabetic codes 121
 - list 122
 - types printed by the disk sort program 24
- MODIFY keyword
 - response in BUILD cycle 98
 - response in CALL cycle 99
 - response in LOAD cycle 94
- multiple input files (Model 15 only) 149
- multiple input support 10
- nine-track tape considerations 16
- nonverify option (column 34--header specifications) 24
 - timing considerations 110
- normal control field
 - changed by ALTSEQ statements 108
 - definition 39
 - how disk sort uses 39
- number of records
 - effect on order of the sequence specifications 8
 - effect on sort time 110
- number of sequence specifications, effect on sort time 111
- numeric constants, coding negative constants 35
- numeric data
 - data formats and column 8 entry combinations 31
 - defined by column 8 record type entry 32
 - definition of types 30
 - numeric data (continued)
 - explanation of positive and negative (signed) numeric data 32
 - indicated by column 8 of record type specifications 31
 - packed numbers 32
 - unpacked numbers 32
- OCL cycles for Model 6
 - functions 91
 - list 91
 - summary of keywords and responses 92
- OCL statements
 - interrelationship with sequence specifications 101
 - supplying for Model 6 16, 91
 - supplying for Model 10 16, 102
 - supplying for Models 12 and 15 17, 102
 - what they tell about disk sort 9
 - what they tell about the input file 9
 - what they tell about the output file 9
 - what they tell about the sequence specifications 9
 - what they tell about the work file 9
 - with conversational OCL 16
 - with online data recorder 16
 - without online data recorder 16
- omit line
 - to exclude records from your sort 27
 - when and how to use 27
- omit set
 - definition 28
 - mixing with include set 29
 - rules for using 28
 - types 29
 - when to use 29
- opposite control field
 - changed by ALTSEQ statements 108
 - column 8 field description entry when sorting different record types 41
 - definition 39
 - how disk sort uses 38
 - when to use 42
- OR lines
 - definition 29
 - describing different record types 28
 - in include sets 29
 - in omit sets 29
- order of sequence specifications
 - affected by number and format of records 8
 - normal order 8
- out-of-sequence lines
 - how to insert in field description specifications 38
 - how to insert in record type specifications 27
- output
 - addrout sort 5
 - OCL statements and error messages 16
 - summary tag along sort 6
 - tag along sort 5

- output file
 - calculating storage tracks for 116
 - considerations 14
 - defined by BUILD cycle 97
 - defined by LOAD cycle 93
 - how position on disk affects sort time 111
 - OCL statements needed if input file is temporary 9
 - overlying input file (RETAIN-S vs RETAIN-T) 13
- output option for SORTR (column 28—header specification) to drop control fields in tag along (SORTR) sorts 24
- output record length for SORTR (columns 29-32—header specifications)
 - calculating when dropping control fields 24
 - calculating when not dropping control fields 24
- overflow field length entry (columns 20-22—field description specifications) 52
- overflow indicator field
 - how to define (columns 7-8—field description specifications) 46
 - how to initialize (column 18—field description specifications) 52
 - when to use 52
- overview of disk sort 3

- packed numeric data
 - number of bits in each digit 32
 - position of sign code 32
 - possible positive and negative sign codes 32
- packed, unpacked, opposite character, or opposite zone control fields 41
- page number (columns 1-2) with line number (columns 3-5) for sequencing (see line number)
- print option (column 27—header specifications), information disk sort can print 24
- programming considerations
 - card input 97
 - data station input 97
 - moving a character 97
 - using disk sort 18

- record length (columns 40-43—header specification for CCP/disk sort), to specify 25, 151
- record length code (column 39—header specification for CCP/disk sort), to specify value 25, 151
- record size
 - calculating for tag along (SORTR) sort 24
 - effect on sort time 110
- record type specifications (disk sort and CCP/disk sort)
 - column summary 26
 - column 6 (line type) 27
 - column 7 (continuation or comments) 30
 - column 8 (C/Z/D/P/U) 30
 - column 19 (field or constant) 35
 - columns you must fill in 26
 - columns 1-2 (page number) 27
- record type specifications (disk sort and CCP/disk sort) (continued)
 - columns 3-5 (line number) 27
 - columns 9-16 (factor 1 location) 34
 - columns 17-18 (relation) 34
 - columns 20-27 (factor 2 field) 35
 - columns 20-39 (factor 2 constant) 35
 - columns 40-72 (comments) 36
 - explanation of column entries 26
 - what they tell disk sort 26
 - when not required 26
- record types
 - described in include set 28
 - described in omit sets 29
- records, how number of affects sort time 110
- relation (columns 17-18—record type specifications)
 - effect on omitting or including records 34
 - to specify factor 1 and factor 2 comparison 34
- replacements for more than one character (column 19—field description specifications), when and how to indicate 46

- sample disk sort timings 114
- sample jobs
 - using addrout (SORTA) and tag along (SORTR) sorts 55
 - using signed control fields 72
 - using summary (SORTRS) sort 67
- sample program
 - instructions for executing on Model 6 146
 - instructions for executing on Models 8, 10, 12, and 15 147
- sequence
 - changing collating sequence of one character 108
 - changing collating sequence of several characters 108
 - standard and alternate collating 107
- sequence (column 18—header specifications), to indicate record sequence 23
- sequence error message
 - field description specifications 38
 - header specifications 23
 - record type specifications 27
- sequencing
 - on field description specifications 38
 - on header specifications 23
 - on record type specifications 27
 - using control fields 42
- sequence specifications
 - for a summary tag along sort (SORTRS) 21
 - for a tag along sort (SORTR) 19
 - for an addrout sort (SORTA) 20
 - how number of affects sort time 111
 - how to enter during BUILD cycle 98
 - how to enter during CALL cycle 99
 - how to enter during LOAD cycle 94
 - how to enter when in KSE-created source library entry 94
 - how to supply 103
 - inserting new ones in field description lines 38
 - inserting new ones in record type lines 27

- sequence specifications (continued)
 - interrelationship with conversational OCL 101
 - maximum number 86
 - normal order 8
 - recommended number for sort job 111
 - using additional to decrease sort time 111
 - ways to supply for Model 6 101
- sequence specification sheet
 - location and purpose of specification types 7
 - types of specifications 7
- set
 - include set 28
 - mixing include and omit sets 29
 - omit set 29
 - rules for using include and omit sets 28
 - types of include sets 28
 - types of omit sets 29
- seven-track tape considerations 15
- severe error message (S) 121
- sign code
 - packed and unpacked numeric data 32
 - placement in packed and unpacked numeric data 33
- signed numbers
 - how to code when used as factor 2 (columns 20-39—record type specifications) 35
 - possible sign codes 32
- size of main storage, effect on sort time 110
- sizes of records, effect on sort time 110
- sort jobs
 - examples of addrout and tag along sorts 55
 - examples of summary sorts 67
 - information you must supply with OCL 9
 - types 5
- sort time
 - affected by location of files 111
 - affected by main storage size 110
 - affected by number of records 110
 - affected by number of sequence specifications 111
 - affected by order of control fields 113
 - affected by record size 110
 - affected by using alternate collating sequence 111
 - affected by using automatic work file allocation 113
- SORTA (addrout sort)
 - contents of sorted records 5
 - example 58
 - sequence specifications 20
 - using summary data fields 40
- sorted file, contents 6
- SORTR (tag along sort)
 - calculating output record length when dropping control fields 23
 - calculating output record length when not dropping control fields 23
 - contents of sorted records 5
 - function 5
 - sequence specifications 19
 - using summary data fields 40
- SORTRS (summary sort)
 - contents of sorted records 6
 - examples 67
 - sequence specifications 21
 - using overflow indicator fields 52
 - using summary data fields 40
- source library member 103
- standard collating sequence
 - specified on header line (column 26) 23
 - table 106
 - when digit portion is used 107
 - when zone and digit portions are used 107
 - when zone portion is used 107
- storage, main
 - amount used by sequence specifications 111
 - how size affects sort time 110
- summary data fields
 - how overflow indicated 52
 - how to define 40
 - maximum number defined 40
 - used by each type of sort 40
- summary overflow field
 - how to define 46
 - how to initialize 52
- summary sort (SORTRS) (see SORTRS)
- summary sort overflow 52
- summary tag along sort (see SORTRS)
- supplying sequence specifications for Model 6
 - advantages and disadvantages of methods 100
- System/3 data
 - alphameric 31
 - how characters represented with the computer 31
 - how signed numeric constants represent on record type specifications 35
 - how to interpret 31
 - in compare operations 30
 - numeric data 31
 - zone and digit portion of characters 31
- tag along and addrout sorts, comparison 14
- tag along sort (SORTR) (see SORTR)
- temporary file, additional OCL statements needed 9
- terminal message (T) 110
- time to sort a file
 - addrout and tag along sort comparison 14
 - affected by alternate collating sequence 111
 - affected by file location on disk 111
 - affected by main storage assigned to disk sort 110
 - affected by number of input records 110
 - affected by number of sequence specifications 111
 - affected by order of control fields 113
 - affected by size of records 110
- timing
 - factors affecting sort time 17
 - samples 114
- to entry (columns 13-16—field description specifications)
 - describing one character field 45
 - to indicate last position of field in input record 44
- to entry (columns 24-27—record type), to identify last position of factor 1 in input record 35
- translation table and alternate collating sequence sheet 109

- unconditional force
 - example 47
 - explanation 39
- unconditional force character, how to define 46
- unpacked numeric data
 - number of bits in each digit 32
 - position of the sign code 33
 - possible positive and negative sign codes 32
- using duplicate control fields 40

- warning message (W) 110
- work file
 - calculating storage tracks for 116
 - considerations 13
 - considering size when programming 18
 - defined during BUILD cycle 96
 - defined during LOAD cycle 93
 - how position on disk affects sort time 111
 - OCL statements needed to describe if temporary 9
- work record length
 - affected by type of output 116
 - affected by type of sort 116
- work record, how used by disk sort 4

- zone, portion of System/3 character 31

- 7-track tape considerations 15
- 9-track tape considerations 16



International Business Machines Corporation

**General Systems Division
4111 Northside Parkway N.W.
P.O. Box 2150
Atlanta, Georgia 30301
(U.S.A. only)**

**General Business Group/International
44 South Broadway
White Plains, New York 10601
U.S.A.
(International)**