Systems

DOS/VSE
Macro Reference

IBM

Preface

This publication contains quick reference information on the data management and system control macros for the experienced programmer and systems support personnel. For the most part, restrictions and programming details have been omitted in order to provide rapid access to the information in this publication. If the information included herein is not sufficient for your purposes, refer to the books in the following publications list.

DOS/VSE Data Management Concepts, GC24-5138

DOS/VSE Macro User's Guide, GC24-5139

DOS/VSE System Management Guide, GC33-5371

OS/VS - DOS/VSE - VM/370 Assembler Language, GC33-4010

Introduction to DOS/VSE, GC33-5370

DOS/VSE System Control Statements, GC33-5376

DOS/VSE Serviceability Aids and Debugging Procedures, GC33-5380

This publication is divided into five chapters. The first chapter contains a list of the macros described later in the book showing their formats and listing the operands permitted for each. The list also contains a succinct description of each macro's function and the number of the page where the detailed de-

scription of the macro may be found. The macros are listed in alphabetical order within the chapter, which serves as an index for the remainder of the book.

Chapter 2 contains an explanation of the macro notation used in this book.

Chapter 3 contains the descriptions of the declarative macros; Chapter 4 contains the imperative macros; Chapter 5 contains the system control macros. These chapters describe the macros in detail. With the following exceptions, the macros are arranged in alphabetic order within each chapter:

- The logic module generation (xxMOD) macros follow the DTFxx macros with which each is associated; for instance, PRMOD follows immediately after DTFPR.

- The line type macros, DFR and DLINT, follow immediately after DTFDR and DRMOD, with which they are associated.

Some of the information included in this book applies only to systems on which VSE/Advanced Functions is installed. In those cases, the VSE/Advanced Functions material is lightly shaded for easy identification.

Chapter 1: Data Management and System Control Macro Formats

Name	Operation	Operands	Use	Page No.
[name]	ATTACH	{entrypoint|(S,entrypoint)|(r1)} ,SAVE={savearea|(S,savearea)|(r2)} [,ECB={ecbname|(S,ecbname)|(r3)}] [,ABSAVE={savearea|(S,savearea)|(r4)}] [,MFG={area|(S,area)|(r5)}]	Initiates a subtask	5-1
[name]	CALL	{entrypoint|(15)} [,(parameterlist)]	Passes control to a specified entry point in another program	5-2
[name]	CANCEL	[ALL]	Terminates a task or sub-task	5-2
[name]	CCB	SYSnnn ,command-list-name [,X'nnnn'] [,senseaddress]	Defines an IOCS command control block	4-1
[name]	CDLOAD	{phasename|(1)} [,PAGE=YES]	Loads a specified phase into the partition GETVIS area	5-3
[name]	CDMOD	[CONTROL=YES] [,CRDERR=RETRY] [,CTLCHR={ASA|YES}] [,DEVICE=nnnn] [,FUNC={R|P|I|RP|RW|RPW|PW}] [,IOAREA2=YES] [,RDONLY=YES] [,RECFORM={FIXUNB|VARUNB|UNDEF}] [,SEPASMB=YES] [,TYPEFLE={INPUT|OUTPUT|CMBND}] [,WORKA=YES]	Defines a logic module for a card reader file	3-6
[name]	CHAP		Lowers the priority of the issuing subtask	5-3
[name]	CHECK	{filename|(1)} [,control-addr|(0)]	Prevents processing until I/O data transfer is complete	4-5
[name]	CHKPT	SYSnnn ,{restart-addr|(r1)} [,end-addr|,(r2)] [,tpointer|(r3)] [,dpointer|(r4)] [,filename|(r5)]	Records the status of your program for later restarting	5-3
[name]	{CLOSE| CLOSER}	{filename1|(r1)} [,filename2|,(r2)]...	Deactivates a file	4-8
[name]	CNTRL	{filename|(1)} ,code [,n1][,n2]	Provides non-data device commands	4-8
[name]	COMRG	[REG=nn]	Places the partition's communication region address in register nn.	5-4
[name]	CPCLOSE	[arglist|(r1)]	Issues a CPCLOSE command to VM/370 to release a print or punch file for output	5-4
[name]	{DAMOD| DAMODV}	[AFTER=YES] [,ERREXT=YES] [,FEOVD=YES] [,HOLD=YES] [,IDLOC=YES] [,RDONLY=YES] [,RECFORM=xxxxxx] [,RELTRK=YES [,RPS=SVA] [,SEPASMB=YES]	Defines a logic module for a direct access file	3-15
[name]	DEQ	{rcbname|(0)}	Releases an ENQ'ed resource	5-4
[name]	DETACH	[SAVE={savearea|(1)}]	Terminates (normally) a subtask	5-5

Name	Operation	Operands	Use	Page No.
[name]	DFR	FONT=xxxx [,REJECT=YES] [,ERASE=YES] [,CHRSET=n] [,EDCHAR=(x,...)] [,BCH=n] [,BCHSER=n] [,NATNHP=YES]	Defines attributes common to a group of line types	3-25
[name]	DIMOD	[IOAREA2=YES] [,RDONLY=YES] [,RPS=SVA] [,SEPASMB=YES] [,TRC=YES] [,TYPEFLE={OUTPUT|INPUT}]	Defines a logic module for a device-independent file	3-20
[name]	DISEN	{filename|(1)}	Stops feeding documents through MICR or OCR devices	4-10
[name]	DLINT	LFR=nn,LINBEG=nn [,IMAGE=YES] [,NOSCAN=(n,n)] [,FLDn=(n,n,NCRIT,xxx)] [,EDITn=(xxxxxx,EDCHAR)] [,FREND=YES]	Describes line types, fields in the line	3-27
[name]	DRMOD	[DEVICE=3886] [,RDONLY=YES] [,SEPASMB=YES] [,SETDEV=YES]	Defines logic modules for a 3886 file	3-24
[name]	DSPLY	{filename|(1)} ,(r2),(r3)	Displays document field on 1287 display scope	4-10
[name]	DTFCD	DEVADDR=SYSxxx ,IOAREA=xxxxxxxx [,ASOCFLE=xxxxxxx] [,BLKSIZE=nnn] [,CONTROL=YES] [,CRDERR=RETRY] [,CTLCHR=xxx] [,DEVICE=nnnn] [,EOFADDR=xxxxxxxx] [,ERROPT=xxxxxx] [,FUNC=xxx] [,IOAREA2=xxxxxxx] [,IOREG=(nn)] [,MODE=xx] [,MODNAME=xxxxxxxx] [,OUBLKSZ=nn] [,RDONLY=YES] [,RECFORM=xxxxxx] [,RECSIZE=(nn)] [,SEPASMB=YES] [,SSELECT=n] [,TYPEFLE=xxxxxx] [,WORKA=YES]	Defines a card or 3881 file	3-1
[name]	DTFCN	DEVADDR=SYSxxx ,IOAREA1=xxxxxxxx [,BLKSIZE=nnn] [,INPSIZE=nnn] [,MODNAME=xxxxxxx] [,RECFORM=xxxxxx] [,RECSIZE=(nn)] [,TYPEFLE=xxxxxx] [,WORKA=YES]	Defines a console file	3-8
[name]	DTFDA	BLKSIZE=nnnn ,DEVICE=nnnn ,ERRBYTE=xxxxxxxx ,IOAREA1=xxxxxxxx	Defines a direct access file	3-10

Name	Operation	Operands	Use	Page No.
		,SEEKADR=xxxxxxxx		
		,TYPEFLE=xxxxxx		
		[,AFTER=YES]		
		[,CONTROL=YES]		
		[,DEVADDR=SYSnnn]		
		[,ERREXT=YES]		
		[,FEOVD=YES]		
		[,HOLD=YES]		
		[,DSKXTNT=n]		
		[,IDLOC=xxxxxxxx]		
		[,KEYARG=xxxxxxx]		
		[,KEYLEN=nnn]		
		[,LABADDR=xxxxxxxx]		
		[,MODNAME=xxxxxxx]		
		[,RDONLY=YES]		
		[,READID=YES]		
		[,READKEY=YES]		
		[,RECFORM=xxxxxx]		
		[,RECSIZE=(nn)]		
		[,RELTYPE=xxx]		
		[,SEPASMB=YES]		
		[,SRCHM=YES]		
		[,TRLBL=YES]		
		[,VERIFY=YES]		
		[,WRITEID=YES]		
		[,WRITEKY=YES]		
		[,XTNTXIT=xxxxxxxx]		
[name]	DTFDI	DEVADDR=SYSxxx	Defines a device-independent file	3-17
		,IOAREA1=xxxxxxxx		
		[,CISIZE=n]		
		[,EOFADDR=xxxxxxxx]		
		[,ERROPT=xxxxxxxx]		
		[,FBA=YES]		
		[,IOAREA2=xxxxxxxx]		
		[,IOREG=(nn)]		
		[,MODNAME=xxxxxxxx]		
		[,RDONLY=YES]		
		[,RECSIZE=nnn]		
		[,SEPASMB=YES]		
		[,TRC=YES]		
		[,WLRERR=xxxxxxxx]		
[name]	DTFDR	DEVADDR=SYSxxx	Defines a 3886 OCR file	3-21
		,FRNAME=xxxxxxxx		
		,FRSIZE=nn		
		,EXITIND=xxxxxxxx		
		,IOAREA1=xxxxxxxx		
		,HEADER=xxxxxxxx		
		,EOFADDR=xxxxxxxx		
		,COREXIT=xxxxxxxx		
		[,DEVICE=3886]		
		[,RDONLY=YES]		
		[,MODNAME=xxxxxxx]		
		[,BLKSIZE=nnn]		
		[,SEPASMB=YES]		
		[,SETDEV=YES]		

Name	Operation	Operands	Use	Page No.
		[,LENCHK=YES]		
		[,MODNAME=xxxxxxxx]		
		[,NOTEPNT=xxxxxx]		
		[,RDONLY=YES]		
		[,READ=xxxxxxx]		
		[,RECFORM=xxxxxx]		
		[,RECSIZE=nnnn]		
		[,REWIND=xxxxxx]		
		[,SEPASMB=YES]		
		[,TPMARK=[YES|NO]		
		[,TYPEFLE=xxxxxx]		
		[,VARBLD=(nn)]		
		[,WLRERR=xxxxxxxx]		
		[,WORKA=YES]		
[name]	DTFOR	COREXIT=xxxxxxxx	Defines a 1287 or 1288 optical reader file	3-53
		,DEVADDR=SYSnnn		
		,EOFADDR=xxxxxxxx		
		,IOAREA1=xxxxxxxx		
		[,BLKFAC=nn]		
		[,BLKSIZE=nn]		
		[,CONTROL=YES]		
		[,DEVICE=xxxxx]		
		[,HEADER=YES]		
		[,HPRMTY=YES]		
		[,IOAREA2=xxxxxxxx]		
		[,IOREG=(nn)]		
		[,MODNAME=xxxxxxxx]		
		[,RECFORM=xxxxxx]		
		[,RECSIZE=(nn)]		
		[,SEPASMB=YES]		
		[,WORKA=YES]		
[name]	DTFPH	TYPEFLE=xxxxxx	Defines a Physical IOCS file	3-58
		[,ASCII=YES]		
		[,CISIZE=n]		
		[,CCWADDR=xxxxxxxx]		
		[,DEVICE=xxxx]		
		[,DEVADDR=SYSxxx]		
		[,HDRINFO=YES]		
		[,LABADDR=xxxxxxxx]		
		[,MOUNTED=xxxxxx]		
		[,XTNTXIT=xxxxxxxx]		
[name]	DTFPR	DEVADDR=SYSxxx	Defines a printer file	3-61
		,IOAREA1=xxxxxxxx		
		[,ASOCFLE=xxxxxxxx]		
		[,BLKSIZE=nnn]		
		[,CONTROL=YES]		
		[,CTLCHR=xxx]		
		[,DEVICE=nnn]		
		[,ERROPT=xxxxxxxx]		
		[,FUNC=xxxx]		
		[,IOAREA2=xxxxxxxx]		
		[,IOREG=(nn)]		
		[,MODNAME=xxxxxxxx]		
		[,PRINTOV=YES]		
		[,RDONLY=YES]		
		[,RECFORM=xxxxxx]		
		[,RECSIZE=(nn)]		
		[,SEPASMB=YES]		
		[,STLIST=YES]		
		[,TRC=YES]		
		[,UCS=xxx]		
		[,WORKA=YES]		

Name	Operation	Operands	Use	Page No.
[name]	FCEPGOUT	{{listname|(1)} |beginaddr,endaddr| [beginaddr,endaddr]}	Forces an area to be paged out	5-7
[name]	FEOV	{filename|(1)}	Forces end-of-volume for magnetic tape file	4-11
[name]	FEOVD	{filename|(1)}	Forces end-of-volume for DASD file	4-11
[name]	FETCH	{phasename|(S,addr)|(1)} [,entrypoint|(S,entrypoint|(0)] [,LIST={listname|(S,listname)|(r1)}] [,SYS=YES] [,DE=YES] [,MFG={area|(S,area)|(r2)}]	Loads a phase; transfers control to it	5-8
[name]	FREE	{filename|(1)}	Makes available to other tasks a previously held track or CI	5-8
[name]	FREEVIS	[ADDRESS={name1|(1)}] [,LENGTH={name2|(0)}] [,SVA=YES]	Releases blocks of virtual storage previously obtained by a GETVIS	5-9
[name]	GENDTL	[ADDR={name1|(S,name1)|(r1)}] [,NAME={name2|(S,name2)|(r2)}] [,CONTROL={E|S}] [,LOCKOPT={1|2}] [,KEEP={NO|YES}] [,OWNER={TASK|PARTITION}] [,LENGTH={NO|YES}]	Generates a DTL (Define The Lock) control block at executed time	5-9
[name]	GENIORB	[ADDRESS={name1|(S,name1)|(1)}] [,LENGTH=fieldlength] ,CCW={name2|(S,name2)|(r2)} ,{DEVICE=SYSxxx| LOGUNIT={name3| (S,name3)|(r3)}} [,FIXLIST={name4|(S,name4)|(r4)}] [,FIXFLAG=(option11,...)] [,IOFLAG=(option21,...)] [,ERREXT={name5|(S,name5)|(r5)}] [,ECB={name6|(S,name6)|(r6)}]	Generates an I/O Request Block at execution time	4-12
[name]	GENL	phasename1,phasename2,... [{,ADDRESS={area|(S,area)|(r1)} ,LENGTH=number}] [{,ADDRESS={DYN|DYNAMIC [,ERREXT={addr|(S,addr)|(r2)}]}]	Generates a local directory list in the partition	5-9
[name]	GET	{filename|(1)} [,workname|,(0)]	Obtains the next sequential logical record from an input file	4-12
[name]	GETIME	[STANDARD|BINARY|TU] [,LOCAL|GMT] [,MFG={area|(S,area)|(r)}]	Obtains the time of day	5-10
[name]	GETVIS	[ADDRESS={name1|(1)}] [,LENGTH={name2|(0)}] [,PAGE=YES] [,POOL=YES] [,SVA=YES]	Obtains a block of virtual storage from a GETVIS area	5-10
[name]	IORB	DSECT=YES or CCW=name1,DEVICE=SYSxxx [,ECB=name2] [,FIXLIST=name3] [,FIXFLAG=(option 11,...)] [,IOFLAG=(option 21,...)]	Generates an I/O Request Block at assembly time	4-13
[name]	ISMOD	[CORDATA=YES] [,CORINDX=YES] [,ERREXT=YES] [,HOLD=YES] [,IOAREA2=YES] ,IOROUT=LOAD|ADD|RETRVE|ADDRTR	Defines a logic module for an indexed sequential file	3-41

Name	Operation	Operands	Use	Page No.
		[,RDONLY=YES] [,RECFORM=FIXUNB|FIXBLK|BOTH] [,RPS=SVA] [,SEPASMB=YES] [,TYPEFLE=RANDOM|SEQNTL|RANSEQ]		
[name]	JDUMP		Produces a hexadecimal dump; teminates the main or subtask	5-11
[name]	JOBCOM	FUNCT={PUTCOM|GETCOM}, AREA={address|(r1)}, LENGTH={length|(r2)}	Permits communication between jobs or job steps in a partition	5-11
[name]	LBRET	{1|2|3}	Returns control to IOCS after label-processing	4-14
[name]	LFCB	SYSxxx,phasename [,NULMSG] [,FORMS=xxxx] [,LPI=n]	Loads the forms control buffer	5-12
[name]	LITE	{filename|(1)} [,light-switches|,(0)]	Lights pocket lamps on 1419 or 1275	4-15
[name]	LOAD	{phasename|(S,address)|(1)} [,loadpoint|(S,loadpoint)|(0)] [,LIST={listname|(S,listname)|(r1)] [,SYS=YES] [,DE=YES] [,TXT=NO] [,MFG={area|(S,area)|(r2)}]	Loads specified phase; returns control to calling phase	5-13
[name]	LOCK	{name|(S,name)|(r)} [,FAIL={RETURN|WAITC|WAIT}]	Enques a resource access request with protection against disallowed	5-14
[name]	MODDTL	ADDR={name1|(S,name1)|(r1)} [,NAME={name2|(S,name2)|(r2)}] [,CONTROL={E|S}] [,LOCKOPT={1|2}] [,KEEP={NO|YES}] [,OWNER={TASK|PARTITION}] [,LENGTH={NO|YES}]	Modifies a DTL (Define The Lock) control block	5-15
[name]	MRMOD	[ADDRESS={SINGLE|DUAL}] [,BUFFERS=nnn] [,SEPASMB=YES]	Defines a logic module for a MICR or OCR file	3-45
[name]	MTMOD	[ASCII=YES] [,CKPTREC=YES] [,ERREXT=YES] [,ERROPT=YES] [,NOTEPNT={YES|POINTS}] [,RDONLY=YES] [,READ={FORWARD|BACK}] [,RECFORM={FIXUNB|FIXBLK|VARUNB| VARBLK|SPNBLK|SPNUNB|UNDEF}]	Defines a logic module for a magnetic tape file	3-51
[name]	MVCOM	to,length,{from|(0)}	Modifies communication region	5-15
[name]	NOTE	{filename|(1)}	Obtains identification for a physical record or logical block	4-15
[name]	{OPEN| OPENR}	{filename1|(r1)} [,filename2|,(r2)],...	Activates a file	4-15
[name]	ORMOD	[BLKFAC=YES] [,CONTROL=YES] ,DEVICE={1287D|1287T} [,IOAREA2=YES] [,RECFORM={FIXUNB|FIXBLK|UNDEF}] [,SEPASMB=YES] [,WORKA=YES]	Defines a logic module for a 1287 or 1288 optical reader file	3-57
[name]	PAGEIN	{{listname|(1)}|beginaddr.endaddr| [,beginaddr.endaddr],...} [,ECB={ecbname|(0)}]	Brings specified areas into real storage	5-16

Name	Operation	Operands	Use	Page No.
[name]	PDUMP	{,address|(r1)}, {address2|(r2)} [,MFG={area|(S,area)|(r3)}]	Produces snapshot hexadecimal dump; processing continues at next instruction	5-16
[name]	PFIX	{{listname|(1)} |beginaddr,endaddr [,beginaddr,endaddr],...}	Brings pages into real storage; fixes them	5-17
[name]	PFREE	{{listname|(1)} |beginaddr,endaddr [,beginaddr,endaddr],...}	Decrements a page's PFIX counter by 1	5-17
[name]	POINTR	{filename|(1)} ,{address|(0)}	Repositions a file to a specified record	4-16
[name]	POINTS	{filename|(1)}	Repositions a file to its begining	4-16
[name]	POINTW	{filename|(1)} ,{address|(0)}	Repositions a file to a specified record	4-16
[name]	POST	{ecbname|(1)} [,SAVE= {savearea|(0)}]	Posts an ECB to a waiting task from the wait state	5-18
[name]	PRMOD	[CONTROL=YES] [,CTLCHR=YES] [,DEVICE=xxxxx] [,ERROPT=YES] [,FUNC=xxxxxx] [,IOAREA2=YES] [,PRINTOV=YES] [,RDONLY=YES] [,RECFORM=xxxxxx] [,SEPASMB=YES] [,STLIST=YES] [,TRC=YES] [,WORK=YES]	Defines a logic module for a printer file	3-65
[name]	PRTOV	{filename|(1)}, {9|12} [,routinename|,(0)]	Specifies printer action when carriage overflow occurs	4-17
[name]	PTMOD	[DEVICE=nnnn] [,RECFORM=xxxxxx] [,SCAN=YES] [SEPASMB=YES] [,TRANS=YES]	Defines a logic module for a paper tape file	3-71
[name]	PUT	{filename|(1)} [,workname|,(0)] [,STLSP= {controlfield|(r)}] [,STLSK= {controlfield|(r)}]	Moves (outputs) a logical record to I/O device	4-17
[name]	PUTR	{filename|(1)} [, {workname1|(0)} ,{workname2|(2)}]	Sends message to operator's console, requiring a reply	4-18
[name]	RCB		Generates a Resource Control Block	5-18
[name]	RDLINE	{filename|(1)}	Reads a 1287 journal tape line in correction mode	4-18
[name]	READ	{filename|(1)} {,KEY|,ID|,MR| ,OR, {name|(r)} | ,DR, {name|(r)|nn,nn} | ,SQ, {area|(0)} [,length|,(r)|,S]}	Transfers data from an input file to an area in virtual storage	4-19
[name]	REALAD	{address|(1)}	Returns a real storage address corresponding to a virtual address	5-18
[name]	RELEASE	(SYSnnn,SYSnnn,...) [,savearea]	Releases programmer logical units	5-19
[name]	RELPAG	{{listname|(1)} |beginaddr,endaddr [,beginaddr,endaddr],...}	Releases specified storage areas	5-19
[name]	RELSE	{filename|(1)}	Skip the remaining records in a block	4-19
[name]	RESCN	{filename|(1)} ,(r1),(r2) ,[n1][,n2]	Rescans a field on an OCR document	4-19
[name]	RETURN	(r1[,r2])	Restores registers, returns control to calling program	5-20

Name	Operation	Operands	Use	Page No.
		[,XECBADR={xecbfield|(S,xecbfield)|(r1)}] [,ACCESS={XPOST|XWAIT}] [,MFG={area|(S,area)|(r2)}]		
[name]	XPOST	XECB={xecbname|(1)} ,POINTREG=(14)	Posts a specified XECB	5-28
[name]	XWAIT	XECB={xecbname|(1)} ,POINTREG=(14)	Waits for a specified XECB to be posted	5-29

Chapter 2: Macro Notation

Macro Fields

Macros, like assembler statements, have a name field, operation field, and operand field. Comments can also be included as in assembler statements, although certain macros require a comment to be preceeded by a comma if the macro is issued without an operand. These macros are; CANCEL, DETACH, FREEVIS, GETIME, GETVIS, TESTF, and TTIMER.

The *name field* in a macro may contain a symbolic name. Some macros (for example, CCB, TECB, or DTFxx) require a name.

The *operation field* must contain the mnemonic operation code of the macro.

The operands in the *operand field* must be written in either positional, keyword, or mixed formats.

Positional Operands

In this format, the parameter values must be in the exact order shown in the individual macro discussion. Each operand, except the last, must be followed by a comma, and no embedded blanks are allowed. If an operand is to be omitted in the macro and following operands are included, a comma must be inserted to indicate the omission. No commas need to be included after the last operand. Column 72 must contain a continuation punch (any nonblank character) if the operands fill the operand field and overflow onto another line.

For example, GET uses the positional format. A GET for a file named CDFILE using WORK as a work area is written:

    GET CDFILE, WORK

Keyword Operands

An operand written in keyword format has this form, for example:

    LABADDR=MYLABELS

where LABADDR is the keyword. MYLABELS is the specification, and LABADDR=MYLABELS is the complete operand.

The keyword operands in the macro may appear in any order, and any that are not required may be omitted. Different keyword operands may be written in the same statement, each followed by a comma except for the last operand of the macro.

Mixed Format

The operand list contains both positional and keyword operands. The keyword operands can be written in any order, but they must be written to the right of any positional operands in the macro.

For additional information on coding macro statements, see *OS/VS -DOS/VSE-VM/370 Assembler Language,* as listed in the Preface.

Notational Conventions

The following conventions are used in this book to illustrate the format of macros:

1. Uppercase letters and punctuation marks (except as described in these conventions) represent information that must be coded exactly as shown.

2. Lowercase letters and terms represent information which you must supply. More specifically, an n indicates a decimal number, an r indicates a decimal register number, and an x indicates an alphameric character.

3. Information contained within brackets [ ] represents an optional parameter that can be included or omitted, depending on the requirements of the program.

4. Stacked options contained within brackets represent alternatives, one of which can be chosen for example:

    $$\left[\begin{array}{c} \text{name} \\ \text{label} \\ \text{address} \end{array}\right]$$
    A name-field symbol in this assembly, or an operand of an EXTRN statement, or * (the location counter).

5. Stacked options contained within braces { } represent alternatives, one of which must be chosen.

6. Items 4 and 5 above may also be shown between brackets and braces, respectively, on one line, that is, unstacked. In that case, the options are separated by OR symbols (|). Examples of this notation are

    {phasename|(1)} [,entrypoint|,(0)]

7. An ellipsis (a series of three periods) indicates that a variable number of items may be included.

8. **filename**   Example of a symbol appearing in the name field of a DTF macro.

9. **n** Self-defining value, such as 3, X'04', (15),B'010'.

10. **length** Absolute expression, as defined in *OS/VS-DOS/VSE-VM/370 Assembler Language*, as listed in the Preface.

11. **{A|B|C}** Underlined elements represent an assumed value in the event an operand is omitted.

12. **(r)** Ordinary register notation. Any register except 0 or 1 to be specified in parentheses.

13. **(0)|(1)** Special register notation (ordinary register notation can be used).

Register Notation

Certain operands can be specified in either of two ways:

1. You may specify the operand directly which results in code that, for example, cannot be executed in the SVA because it is not reentrant.

2. You may load the address of the value into a register before issuing the macro. This way the generated code is reentrant and may be executed in the SVA. When using register notation, the register should contain only the specific address and high order bits should be set to 0.

When the macro is assembled, instructions are generated to pass the information contained in the specified register to IOCS or to the supervisor. For example, if an operand is written as (8), and if the corresponding parameter is to be passed to the supervisor in register 0, the macro expansion contains the instruction LR 0,8. This is an example of ordinary register notation.

You can save both storage and execution time by using what is known as special register notation. In this method, the operand is shown in the format description of the macro as either (0) or (1), for example. This notation is special because the use of registers 0 and 1 is allowed only for the indicated purpose.

If special register notation is indicated by (0) or (1) in a macro format description and you use ordinary register notation, the macro expansion will contain an extra LR instruction.

The format description for each macro shows whether special register notation can be used, and for which operands. The following example indicates that the filename operand can be written as (1) and the workname operand as (0):

GET {filename|(1)} [workname|,(0)]

If either of these special register notations is used, your program must load the designated parameter register before executing the macro expansion. Ordinary register notation can also be used.

Operands in (S,address) Notation

Certain system control macros (for instance, ATTACH, GENIORB, GENL, LOAD) allow three notations for an operand:

1. Register notation, as described in the preceding paragraph.

2. Notation as a relocatable expression which, in the macro expansion, results in an A-type address constant.

3. Notation in the form (S,address). In the macro expansion, an explicit address (that is: an assembler instruction address in base-displacement form) is generated. Address can be specified either as a relocatable expression --for example: (S,RELOC), or as two absolute expressions, the first of which represents the displacement and the second, the base register --for example: (S,512(12)).

You should consider using this notation if your program is to be reenterable. In a reenterable program, macro operands often refer to fields in dynamic storage. The (S,address) format offers an alternative to register notation: if two or more of such operands have to be provided for one macro, there is no need for loading addresses into that many registers.

Declarative Macro Statements

The operands of the DTFxx and the logic module generation macros are written in assembler format statements. The first statement is the header and the continuations following are the detail statements. The header contains:

- The symbolic name of the file in the name field. In a DTF, the symbolic file name may have as many as seven characters. The file name may also be required on standard label job control statements and in certain macros as operands; it must be the same as that used in the DTF header. For a logic module, the name may not be required.

- The mnemonic operation code of the macro in the operation field.

- Keyword operands in the operand field, as required.

- A continuation character in column 72, if required.

Note: Avoid using IJ as the first two letters when defining symbols as they may conflict with IOCS symbols beginning with IJ. Avoid symbols that are identical to a filename plus a single character suffix because IOCS generates symbols by concatenating the filename with an additional character. For the filename RECIN, for example, IOCS generates the symbols RECINS, RECINI, etc.

The details follow the header and may be arranged in any convenient order. Each continuation line must begin in column 16. If more than one operand is written on a detail line, they must be separated by a comma only. Except for the final detail line, there must be a comma immediately following each operand and have a continuation character in column 72. You may include a comment on a header or a detail line if there is room between a space following the last operand on a line and column 72.

Chapter 3: Declarative Macros

DTFCD Macro
This macro defines a file for a card reader.

Applies to					
Input	**Output**	**Combined**			
x	x	x	M	DEVADDR=SYSxxx	Symbolic unit for reader-punch used for this file
x	x	x	M	IOAREA1 =xxxxxxx	Name of first I/O area, or separate input area if TYPEFLE=CMBND and IOAREA2 are specified.
x	x		O	ASOCFLE=xxxxxx	Name for FUNC=RP, RW, RPW, PW
x	x	x	O	BLKSIZE=nnn	Length of one I/O area, in bytes. If omitted, 160 is assumed for a column binary on the 2560,3504,3505, or 3525; 96 is assumed for the 2596 or 5424/5425, otherwise 80 is assumed.
x	x	x	O	CONTROL=YES	CNTRL macro used for this file. Omit CTLCHR for this file. Does not apply to 2501.
	x		O	CRDERR=RETRY	RETRY if punching error is detected. Applies to 2520 and 2540 only.
	x		O	CTLCHR=xxx	(YES or ASA). Data records have control character. YES for S/370 character set; ASA for American National Standards Institute character set. Omit if TYPEFLE=CMBND. Omit CONTROL for this file.
x	x	x	O	DEVICE=nnnn	(1442, 2501, 2520, 2540, 2560P, 2560S, 2596, 3504, 3505, 3525, 5425P, or 5425S). If omitted, 2540 is assumed. Specify 5425P/S for 5424/5425(P/S).
x		x	O	EOFADDR=xxxxxxx	Name of your end-of-file routine.
x	x		O	ERROPT=xxxxxx	IGNORE, SKIP, or name. Applies to 2560, 3504, 3505, 3525 and 5424/5425 only.
x	x		O	FUNC=xxx	R, P, I, RP, RW, RPW, PW. Applies to 2560, 3525, and 5424/5425 only.
x	x	x	O	IOAREA2=xxxxxx	Name of second I/O area, or separate output area if TYPEFLE=CMBND. Not allowed if FUNC=RP, RW, RPW, or PW. Not allowed for output file if ERROPT=IGNORE.
x	x		O	IOREG=(nn)	Register number, if two I/O areas used and GET or PUT does not specify a work area. Omit WORKA.
x	x		O	MODE=xx	(E or C) for 2560. (E, C, O, R, EO, ER, CO, CR) for 3504 and 3505. (E, C, R, ER, CR) for 3535. If omitted, E is assumed.
x	x	x	O	MODNAME=xxxxxxx	Name of CDMOD logic module for this DTF. If omitted, IOCS generates standard name.

M=Mandatory
O=Optional

Figure 3-1. DTFCD macro operands (Part 1 of 2).

Applies to					
Input	Output	Combined			
		x	O	OUBLKSZ=nn	Length of IOAREA2 if TYPEFLE=CMBND. If OUBLKSZ omitted, length specified by BLKSIZE is assumed for IOAREA2.
x	x	x	O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module.
x	x	x	O	RECFORM=xxxxxx	(FIXUNB, UNDEF, or VARUNB). If omitted, FIXUNB is assumed. Input or combined files always FIXUNB.
	x		O	RECSIZE=(nn)	Register number if RECFORM=UNDEF. General registers 2-12, written in parentheses.
x	x	x	O	SEPASMB=YES	DTFCD is to be assembled separately.
x	x		O	SSELECT=n	(1 or 2) for 1442, 2520, 2596, 3504, or 3525. (1, 2, or 3) for 2540. (1, 2, 3, 4, or 5) for 2560. (1, 2, 3, or 4) for 5424/5425. Stacker-select character.
x	x	x	O	TYPEFLE=xxxxxx	(INPUT, OUTPUT, or CMBND) If omitted INPUT assumed. CMBND may be specified for 1442N1, 2520B1, or 2540 punch-feed-read only.
x	x	x	O	WORKA=YES	GET or PUT specifies work area. Omit IOREG. Not allowed for output file if ERROPT=IGNORE.

M=Mandatory
O=Optional

Figure 3-1. DTFCD macro operands (Part 2 of 2).

Code in FUNC= operand	filename specification in ASOCFLE= operand of		
	read DTFCD	punch DTFCD	print DTFPR
FUNC=PW		filename of print DTFPR	filename of punch DTFCD
FUNC=RP	filename of punch DTFCD	filename of read DTFCD	
FUNC=RPW	filename of punch DTFCD	filename of print DTFPR	filename of read DTFCD
FUNC=RW	filename of print DTFPR		filename of read DTFCD

Examples:
1. If FUNC=PW is specified,
   a. specify the filename of the print DTFPR in the ASOCFLE operand of the punch DTFCD and
   b. Specify the filename of the punch DTFCD in the ASOCFLE operand of the print DTFPR.
2. If FUNC=RPW is specified,
   a. specify the filename of the punch DTFCD in the ASOCFLE operand of the read DTFCD, and
   b. specify the filename of the print DTFPR in the ASOCFLE operand of the punch DTFCD, and
   c. specify the filename of the read DTFCD in the ASOCFLE operand of the print DTFPR.

Figure 3-2. ASOCFLE operand usage with print associated files.

ASOCFLE=filename: This operand is used together with the FUNC operand to define associated files for the 2560, 3525, or 5424/5425. (For a description of associated files see the *DOS/VSE Macro User's Guide* , as listed in the Preface.) ASOCFLE specifies the filename of associated read, punch, or print files, and enables macro sequence checking by

the logic module of each associated file. One filename is required per DTF for associated files.

BLKSIZE=n: Enter the length of the I/O area (IOAREA1). If the record format is variable or undefined, enter the length of the largest record. If the operand FUNC=I is specified for the 2560 or 3525, the length specified for BLKSIZE must be 80 data

bytes if CTLCHR=YES or if ASA is not specified, or 81 if CTLCHR=YES or if ASA is specified.

CONTROL=YES: This operand is specified if a CNTRL macro is to be issued for a file. If this operand is specified, CTLCHR must be omitted. The CNTRL macro cannot be used for an input file with two I/O areas (that is, when the IOAREA2 operand is specified).

This operand must not be specified for an input file used in association with a punch file (when the operand FUNC=RP or RPW is specified) on the 2560, 3525, or 5424/5425; in this case, however, this operand can be specified in the DTFCD for the associated punch file.

CRDERR=RETRY: This operand applies to card output on the 2520 or 2540. It specifies the operation to be performed if an error is detected. From this specification, IOCS generates a retry routine and a save area for the card punch record.

If a punching error occurs, it is usually ignored and operation continues. The error card is stacked in stacker P1 (punch), while correct cards are stacked in the stacker you select. If the CRDERR=RETRY operand is included and an error condition occurs, IOCS also notifies the operator and then enters the wait state. The operator can either terminate the job, ignore the error, or instruct IOCS to repunch the card.

CTLCHR={ASA|YES}: This operand is required if first-character control is to be used on an output file. ASA denotes the American National Standards Institute, Inc. character set. YES denotes the EBCDIC character set. *Appendix A* of *DOS/VSE Macro User's Guide* contains a complete list of codes. This entry does not apply to combined files. If this operand is specified, CONTROL must be omitted.

DEVADDR= {SYSIPT|SYSPCH|SYSRDR| SYSnnn}:
This operand specifies the symbolic unit to be associated with a file. The symbolic unit represents an actual I/O device address and is used in the ASSGN job control statement to assign the actual I/O device address to the file.

SYSIPT, SYSPCH, or SYSRDR must not be specified:

- for the 2596
- for the 3881
- for 1442, 2520, or 2540 combined files (TYPEFLE=CMBND)
- for 2560, 3525, or 5424/5425 associated files (FUNC=RP, RW, RPW, or PW)

- if the operand FUNC=I is specified
- if the MODE operand is specified with the C, O, or R parameters.

DEVICE= {2540|1442|2501|2520|2560P| 2560S|2596|3504|3505|3525| 5425P|5425S|3881}:
This operand specifies the I/O device associated with a file. The "P" and "S" included with the "2560" and "5425" parameters specify primary or secondary input hoppers. Specify 5425P/S for 5424P/S.

Note: Modifications to the double buffering concept in LIOCS provide for increased throughput for SYSIPT or SYSRDR files on 2501 card readers attached to a S/370-115-2, S/370-125-0, or S/370-125-2. If you specify DEVICE=2501 together with DEVADDR=SYSIPT|SYSRDR and IOAREA2=name, this will lead to the generation of a second CCB with its CCW pointing to the second I/O area.

EOFADDR=name: This entry must be included for input and combined files and specifies the symbolic name of your end-of-file routine. IOCS automatically branches to this routine on an end-of-file condition. In your routine you can perform any operations required for the end of the file (you generally issue a CLOSE instruction for the file).

IOCS detects end-of-file conditions in the card reader by recognizing the characters /* punched in card columns 1 and 2 (column 3 must be blank). If the system logical units SYSIPT and SYSRDR are assigned to a 5424/5425, IOCS requires that the /* card, indicating end-of-file, be followed by a blank card. An error condition results if cards are allowed to run out without a /* trailer card (and without a /& card to indicate end-of-job).

ERROPT= {IGNORE|SKIP|name}: This operand specifies the error exit option used for an input or output file on a 2560, 3504, 3505, 3525, or 5424/5425. Either IGNORE, SKIP, or the symbolic name of an error routine can be specified for input files. Only IGNORE can be specified for output files. This operand must be omitted when using 2560 or 5424/5425 associated output files. The functions of these parameters are described below.

IGNORE indicates that the error is to be ignored. The address of the record in error is put in register 1 and made available for processing. For output files, byte 3, bit 3 of the CCB is also set on (see Figure 4-2); you can check this bit and take the appropriate action to recover from the error. Only one I/O area and no work area is permitted for output files. When IGNORE is specified for an input file associated with a punch file (FUNC=RP or RPW) and an error occurs, a PUT for the card in error must nevertheless be given for the punch file.

SKIP indicates that the record in error is not to be made available for processing. The next card is read and processing continues.

If name is specified, IOCS branches to your routine when an error occurs, where you may perform whatever actions you desire. Register 1 contains the address of the record in error, and register 14 contains the return address. GET macros must not be issued in the error routine for cards in the same device (or in the same card path for the 2560 or 5424/5425). If the file is an associated file, PUT macros must not be issued in the error routine for cards in the same device (for the 2560 or 5424/ 5425 this applies to cards in either card path). If any other IOCS macros are issued in the routine, register 14 must be saved. If the operand RDONLY=YES is specified, register 13 must also be saved. At the end of your routine, return to IOCS by branching to the address in register 14. If the input file is associated with an output file (FUNC=RP, RPW, or RW), no punching or printing must be done for the card in error. IOCS continues processing by reading the next card.

Note: When ERROPT is specified for an input file and an error occurs, there is a danger that the /* end-of-file card may be lost. This is because IOCS, after taking the action for the card in error specified by the ERROPT operand, returns to normal processing by reading the next card which is assumed to be a data card. If this card is in fact an end-of-file card, the end-of-file condition cannot be recognized.

FUNC= {R|P|I|RP|RW|RPW|PW}: This operand specifies the type of file to be processed by the 2560, 3525, or 5424/5425. R indicates read, P indicates punch, and W indicates print.

When FUNC=I is specified, the file will be both punched and interpreted; no associated file is necessary to achieve this. The information printed will be the same as the information punched, in contrast to FUNC=PW, where any relation between the information printed and the information punched is determined by your program. When FUNC=I is specified the file can have only one I/O area.

RP, RW, RPW, and PW are used, together with the ASOCFLE operand, to specify associated files; when one of these parameters is specified for one file, it must also be specified for the associated file(s). Each of the associated files can have only one I/O area.

IOAREA1=name This operand specifies the name of the input or output area used for this file.

If issued for a combined file, this operand specifies the input area. If IOAREA2 is not specified, the area specified in this operand is used for both input and output.

IOAREA2=name: This operand specifies the name of a second I/O area. If the file is a combined file and the operand is specified, the designated area is an output area.

If this operand is specified for the 3881, the IOREG operand must also be specified.

This operand must not be specified if, for the FUNC operand, any of the parameters I, RP, RPW, RW, or PW is specified or if, for an output file, ERROPT=IGNORE is specified.

IOREG=(r): If work areas are not used but two input or output areas are, this operand specifies the register (any of 2 through 12) in which IOCS puts the address of the record. For output files, IOCS puts the address where the user can build a record. This operand cannot be used for combined files.

This operand must be specified for the 3881 if the IOAREA2 operand is specified.

MODE= {E|C|O|R|EO|ER|CO|CR}: This operand specifies the mode used to process an input or output file for a 2560, 3504, 3505, or 3525. E indicates normal EBCDIC mode; C indicates column binary mode; O indicates optical mark read (OMR) mode; R indicates read column eliminate mode. E is also assumed if only O or R is specified.

For the 2560, only E and C are valid entries.

Valid entries for the 3504 and 3505 are E, C, O, R, EO, ER, CO, and CR. Valid entries for the 3525 are E, C, R, ER, and CR. If O or R is specified (with or without E or C), a format descriptor card defining the card columns to be read, or eliminated, must be provided. See OMR considerations in the *DOS/VSE Macro User's Guide*, as listed in the Preface, for instructions on how to write this card as well as on how to code and process OMR data.

Only E is valid for SYSIPT, SYSPCH, or SYSRDR. O and R (with or without E or C) cannot be specified for output files. E is assumed if the MODE operand is omitted.

MODNAME=name: This operand is used to specify the name of the logic module that will be used with the DTF table to process the file. If the logic module is assembled with the program, MODNAME must specify the same name as the CDMOD macro.

If this operand is omitted, standard names are generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called.

OUBLKSZ=n: This operand is used in conjunction with IOAREA2, but only for a combined file. Enter the maximum number of characters to be transferred at one time. If this entry is not included and IOAREA2 is specified, the same length as defined by BLKSIZE is assumed.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. (In the case of double buffering support for the 2501 Card Reader, the save area must be 76 bytes to include the second CCB generated.) Each task should have its own uniquely defined save area. Each time an imperative macro (except OPEN or OPENR) is issued, register 13 must contain the address of the save area associated with that task. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If an ERROPT routine issues I/O macros using the same read-only module that caused control to pass to the error routine, your program must provide another save area. One save area is used for the normal I/O operations, and the second for I/O operations in the ERROPT routine. Before returning to the module that entered the ERROPT routine, register 13 must contain the save area address originally specified for the task.

If this operand is omitted, the module generated is not reenterable, and no save area is required.

RECFORM= {FIXUNB|VARUNB|UNDEF}: This operand specifies the record format of the file: fixed length, variable length, or undefined. If the record format is fixed unblocked (FIXUNB,) this operand may be omitted. This operand must specify FIXUNB if you also specified one of the following:
    TYPEFLE=INPUT
    TYPEFLE=CMBND,
    FUNC=I
    DEVICE=3881.

RECSIZE=(r) For undefined records, this operand specifies the register (one of 2 through 12) that contains the length of the output record. You must load the length of each record into the specified register before you issue the PUT macro for the record.

SEPASMB=YES: Include this operand only if the DTFCD is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and defines the filename as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

SSELECT=n: This operand specifies the valid stacker-select character for a file. If this entry is not specified, cards are selected into NR (normal read) or NP (normal punch) stackers. For the 5424/5425, cards from hopper 1 are placed in stacker 1 and cards from hopper 2 are placed in stacker 5 (or 4).

This operand must not be specified for combined files, for files on the 3881, for 2560, 3525, or 5424/5425 read files associated with punch files (FUNC=RP or FUNC=RPW); in this case the SSELECT=n operand may be specified for the associated output file. For further information, see "CNTRL Macro."

Note: When this operand is used with a device other than a 1442 or 2596, the program ignores CONTROL=YES with input files.

TYPEFLE= {INPUT|OUTPUT|CMBND}: This operand specifies whether a file is input, output, or combined. A combined file can be specified for a 1442 or 2520 or for a 2540 with the punch-feed-read feature. TYPEFLE=CMBND is applicable if both GETs and PUTs are issued for the same card file.

Only TYPEFLE=INPUT can be specified for the 3881. If OUTPUT or CMBND is specified, the DTF defaults to DEVICE=2540 and a non-executable CDMOD logic module is produced. The MNOTE "Improper device. 2540 assumed." is then printed at assembly time. If the operand is omitted, INPUT is assumed.

WORKA=YES: If I/O records are processed in work areas instead of in the I/O areas, specify this operand. You must set up the work area in storage. The address of the work area, or a general-purpose register which contains the address, must be specified in each GET and PUT macro.

If ERROPT=IGNORE is specified for an output file or if DEVICE=3881, WORKA=YES must not be specified.

CDMOD Macro

Listed here are the operands you can specify for CDMOD. The first card contains CDMOD in the operation field and may contain a module name in the name field.

CONTROL=YES: Include this operand if the CNTRL macro is used with the module and its associated DTFs. The module also processes files for which the CNTRL macro is not used.

If this operand is specified, the CTLCHR operand must not be specified. This operand cannot be specified if IOAREA2 is used for an input file.

This operand must not be specified for an input file used in association with a punch file (when the operand FUNC=RP or RPW is specified) on the 2560, 3525, or 5424/5425; in this case, however, this operand can be specified in the DTFCD and CDMOD for the associated punch file.

CRDERR=RETRY: Include this operand if error retry routines for the 2540 and 2520 punch-equipment check are included in the module. Whenever this operand is specified, any DTF used with the module must also specify the same operand. This operand does not apply to an input or a combined file.

CTLCHR= {ASA|YES}: Include this operand if first character stacker select control is used. Any DTF to be used with this module must have the same operand. If CTLCHR is included, CONTROL must not be specified. This operand does not apply to a combined file or to an input file.

DEVICE= {2540|1442|2501|2520|2560P|
2560S|2596|3504|3505|3525|
5425P|5425S|3881}:
Include this operand to specify the I/O device used by the module. The "P" and "S" included with the "2560" and "5425" parameters specify primary or secondary input hoppers; regardless of which is specified, however, the module generated will handle DTFs specifying either hopper. Specify 5425P/S for 5424P/S.

Any DTF to be used with this module must have the same operand (except as just noted concerning the 'P' and 'S' specification for the 2560 or 5425).

FUNC= {R|P|I|RP|RW|RPW|PW} This operand specifies the type of file to be processed by the 2560, 3525, or 5424/5425. Any DTF used with the module must have the same operand. R indicates read, P indicates punch, and W indicates print.

When FUNC=I is specified, the file will be both punched and interpreted; no associated file is necessary to achieve this.

RP, RW, RPW, and PW specify associated files; when one of these parameters is specified for one file, it must also be specified for the associated file(s). Associated files can have only one I/O area each.

IOAREA2=YES: Include this operand if a second I/O area is used. Any DTF used with the module must also include the IOAREA2 operand. This operand is not required for combined files. This operand is not valid for associated files.

RDONLY=YES: This operand causes a read-only module to be generated. Whenever this operand is specified, any DTF used with the module must have the same operand.

RECFORM= {FIXUNB|VARUNB|UNDEF}:
This operand specifies the record format: fixed-length, variable-length, or undefined. Any DTF used with the module must have the same operand. If TYPEFLE=INPUT, TYPEFLE=CMBND, or FUNC=I, this operand must be FIXUNB. For the 3881, only RECFORM=FIXUNB is valid. If this operand is omitted for the 3881, IOCS assumes RECFORM=FIXUNB.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and defines the module name as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

TYPEFLE= {INPUT|OUTPUT|CMBND}: This operand generates a module for either an input, output, or combined file. Any DTF used with the module must have the same operand. For the 3881, only TYPEFLE=INPUT is valid. If the operand is omitted, INPUT is assumed.

WORKA=YES: This operand must be included if records are to be processed in work areas instead of in I/O areas. Any DTF used with the module must have the same operand. This operand is not valid for the 3881.

Standard CDMOD Names
Each name begins with a 3-character prefix (IJC) and continues with a 5-character field corresponding

to the options permitted in the generation of the module.

CDMOD name = IJCabcde

a = F RECFORM=FIXUNB (always for INPUT, CMBND, or FUNC=I files)
  = V RECFORM=VARUNB
  = U RECFORM=UNDEF

b = A CTLCHR=ASA (not specified if CMBND)
  = Y CTLCHR=YES
  = C CONTROL=YES
  = Z CTLCHR or CONTROL not specified

c = B RDONLY=YES and TYPEFLE=CMBND
  = C TYPEFLE=CMBND
  = H RDONLY=YES and TYPEFLE=INPUT
  = I TYPEFLE=INPUT
  = N RDONLY=YES and TYPEFLE=OUTPUT
  = O TYPEFLE=OUTPUT

d = Z WORKA and IOAREA2 not specified
  = W WORKA=YES
  = I IOAREA2=YES
  = B WORKA and IOAREA2
  = Z WORKA=YES not specified (CMBND file only)

e = 0 DEVICE=2540, 3881
  = 1 DEVICE=1442, 2596
  = 2 DEVICE=2520
  = 3 DEVICE=2501
  = 4 DEVICE=2540 and CRDER
  = 5 DEVICE=2520 and CRDERR
  = 6 DEVICE=3505 or 3504
  = 7 DEVICE=3525 and FUNC=R/P or omitted
  = 8 DEVICE=2560 and FUNC=R/P or omitted
  = 9 DEVICE=5425 and FUNC=R/P or omitted
  = A DEVICE=3525 and FUNC=RP
  = B DEVICE=3525 and FUNC=RW
  = C DEVICE=3525 and FUNC=PW
  = D DEVICE=3525 and FUNC=I
  = E DEVICE=3525 and FUNC=RPW
  = F DEVICE=2560 and FUNC=RP
  = G DEVICE=2560 and FUNC=RW
  = H DEVICE=2560 and FUNC=PW
  = I DEVICE=2560 and FUNC=I
  = J DEVICE=2560 and FUNC=RPW
  = K DEVICE=5425 and FUNC=RP
  = L DEVICE=5425 and FUNC=RW
  = M DEVICE=5425 and FUNC=PW
  = N DEVICE=5425 and FUNC=I
  = O DEVICE=5425 and FUNC=RPW

Subset/Superset CDMOD Names

Figure 3-3 shows the subsetting and supersetting allowed for CDMOD names. All but one of the parameters are exclusive (that is, do not allow supersetting). A module name specifying C (CONTROL) in the b location is a superset of a module name specifying Z (no CONTROL or CTLCHR). A module with the name IJCFCIW0 is a superset of a module with the name IJCFZIW0.

```
          *  *  *  *  *
    I  J  C  F  A  B  B  O
             V  Y  C  I  1
             U  +  H  W  2
                C  I  Z  3
                Z  N     4
                   O     5
                         6
                         7
                         8
                         9
                         A
                         B
                         C
                         .
                         .
                         .
                         M
                         N
                         O
```

+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.
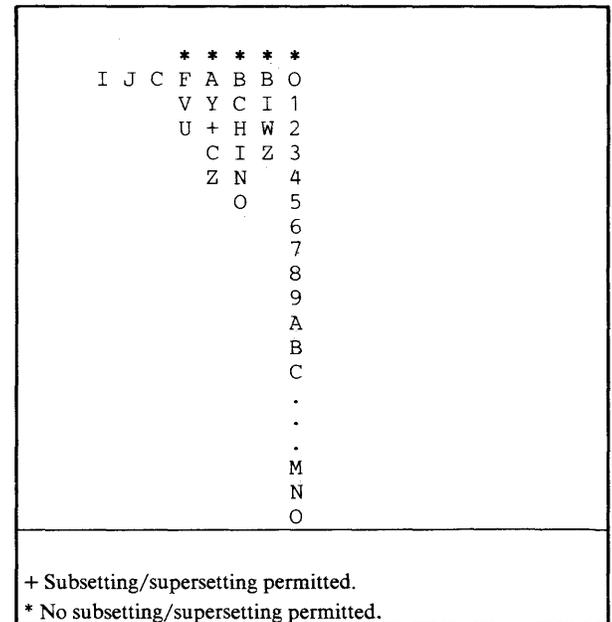
Figure 3-3. Subsetting and supersetting of CDMOD names.

DTFCN Macro

DTFCN defines an input or output file that is processed on a 3210 or 3215 console printer-keyboard, or a display operator console. DTFCN provides GET/PUT logic as well as PUTR logic for a file.

M	DEVADDR = SYSxxx	Symbolic unit for the console used for this file.
M	IOAREA1 =xxxxxxx	Name of I/O area.
O	BLKSIZE=nnn	Length in bytes of I/O area (for PUTR macro usage, length of output part of I/O area). If RECFORM=UNDEF, max. is 256. If omitted, 80 is assumed.
O	INPSIZE = nnn	Length in bytes for input part of I/O area for PUTR macro usage.
O	MODNAME = xxxxxxx	Logic module name for this DTF. If omitted, IOCS generates a standard name. The logic module is generated as part of the DTF.
O	RECFORM=xxxxxx	(FIXUNB or UNDEF). If omitted, FIXUNB is assumed.
O	RECSIZE =(nn)	Register number if RECFORM=UNDEF. General registers 2-12, written in parentheses.
O	TYPEFLE=xxxxxx	(INPUT, OUTPUT, or CMBND). INPUT processes both input and output. CMBND must be specified for PUTR macro usage. If omitted, INPUT is assumed.
O	WORKA=YES	GET or PUT specifies work area.

M = Mandatory
O = Optional

Figure 3-4. DTFCN macro operands.

BLKSIZE=n: This operand specifies the length of the I/O area; if the PUTR macro is used (TYPEFLE=CMBND is specified), this operand specifies the length of the output part of the I/O area. For the undefined record format, BLKSIZE must be as large as the largest record to be processed. The length must not exceed 256 characters.

If the console buffering option is specified at system generation time and the device is assigned to SYSLOG, physical IOCS can increase throughput for each actual output record not exceeding 80 characters. This increase in throughput results from starting the output I/O command and returning to the program before output completion. Regardless of whether or not output records are buffered (queued on an I/O completion basis), they are always printed or displayed in a first-in-first-out (FIFO) order.

DEVADDR= {SYSLOG|SYSnnn}: This operand specifies the symbolic unit associated with the file. In a multiprogramming environment, DEVADDR=SYSLOG must be specified to obtain partition identification prefixes (BG, F1, F2, F3, and F4) for message identification.

DEVADDR=SYSLOG must be specified if your DTFCN macro includes TYPEFLE=CMBND.

INPSIZE=n: This operand specifies the length of the input part of the I/O area for PUTR macro usage.

IOAREA1=name: This operand specifies the name of the I/O area used by the file. For PUTR macro usage, the first part of the I/O area is used for output,

and the second part is used for input. The lengths of these parts are specified by the BLKSIZE and INPSIZE operands respectively. The I/O area is not cleared before or after a message is printed, or when a message is canceled and reentered on the console.

MODNAME=name: This operand specifies the name of the logic module generated by this DTFCN macro. If this entry is omitted, standard module names are generated for the logic module.

A module name must be given when two phases (each containing a DTFCN macro) are link-edited into the same program. Under such conditions, omission of this operand results in unresolved address constants.

RECFORM= {FIXUNB|UNDEF}: This operand specifies the record format of the file: fixed length or undefined. FIXUNB must be specified if TYPEFLE=CMBND is specified. FIXUNB is assumed if the RECFORM operand is omitted.

RECSIZE=(r): For undefined records, this operand is required for output files and is optional for input files. It specifies a general register (2 to 12) that contains the length of the record. On output, you must load the length of each record into the specified register before you issue a PUT macro. If specified for input files, IOCS provides the length of the record transferred to storage.

TYPEFLE={INPUT|**OUTPUT|CMBND**}: This operand specifies a file as input, output, or combined. If INPUT is specified, code is generated for both input and output files. If OUTPUT is specified, code is provided for output files only.

CMBND must be specified if you use the PUTR macro. CMBND specifies that coding be generated for both input and output files; in addition, coding is generated to allow usage of the PUTR macro to en-sure that messages requiring operator action are not deleted from the console. When CMBND is specified, DEVADDR=SYSLOG must also be specified.

WORKA=YES: This operand indicates that a work area is used with the file. A GET or PUT macro moves the record to or from the work area. A PUTR macro moves the record from *and* to the work area.

DTFDA Macro

The DTFDA macro defines a file for Direct Access Method (DAM) processing.

Applies to				
Input	Output			
x	x	M	BLKSIZE=nnnn	Length of one I/O area, in bytes
x	x	M	DEVICE=nnnn	(2311, 2314, 3330, 3340, 3350). If omitted, 2311 is assumed. Specify any device for a 3330-11 or 3350.
x	x	M	ERRBYTE=xxxxxxx	Name of 2-byte field for error/status codes supplied by IOCS
x	x	M	IOAREA1=xxxxxxx	Name of I/O area
x	x	M	SEEKADR=xxxxxxx	Name of track-reference field
x	x	M	TYPEFLE=xxxxxx	(INPUT or OUTPUT)
	x	O	AFTER=YES	WRITE filename, AFTER or WRITE filename, RZERO macro is used for this file
x	x	O	CONTROL=YES	CNTRL macro is used for this file
x	x	O	DEVADDR=SYSnnn	Symbolic unit required only when no extent statement is provided
x	x	O	ERREXT=YES	Nondata transfer errors are to be indicated in ERRBYTE
x		O	FEOVD=YES	Support for sequential disk end of volume records is desired
x		O	HOLD=YES	Employ the track hold function
x	x	O	DSKXTNT=n	Indicates the number (n) of extents for a relative ID
x	x	O	IDLOC=xxxxxxxx	Name of field in which IOCS stores the ID of a record
x	x	O	KEYARG=xxxxxxx	Name of key field if READ filename,KEY; or WRITE filename,KEY; or WRITE filename,AFTER is used for this file
x	x	O	KEYLEN=nnn	Number of bytes in record key if keys are to be processed. If omitted, IOCS assumes zero (no key)
x	x	O	LABADDR=xxxxxxx	Name of your routine to check/write user labels
x	x	O	MODNAME=xxxxxxx	Name of DAMOD logic module for this DTF. If omitted, IOCS generates standard name
x	x	O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module
x		O	READID=YES	READ filename, ID macro is used for this file
x		O	READKEY=YES	READ filename, KEY macro is used for this file
x	x	O	RECFORM=xxxxxx	(FIXUNB, SPNUNB, VARUNB, or UNDEF). If omitted, FIXUNB is assumed
x	x	O	RECSIZE=(nn)	Register number if RECFORM=UNDEF
x	x	O	RELTYPE=xxx	(DEC or HEX). Indicates decimal or hexadecimal relative addressing
x	x	O	SEPASMB=YES	DTFDA is to be assembled separately
x	x	O	SRCHM=YES	Search multiple tracks, if record reference is by key
	x	O	TRLBL=YES	Process trailer labels, LABADDR must be specified
	x	O	VERIFY=YES	Check disk records after they are written.
x	x	O	WRITEID=YES	WRITE filename, ID macro is used for this file
x	x	O	WRITEKY=YES	WRITE filename, KEY macro is used for this file
x	x	O	XTNTXIT=xxxxxxx	Name of your routine to process extent information

M = Mandatory
O = Optional

Figure 3-5. DTFDA macro.

AFTER=YES: This operand must be included if any records (or an additional record) are written in a file by a formatting WRITE (count, key, and data) following the last record previously written on a track. The remainder of the track is erased. That is, whenever either of the macros

WRITE filename,AFTER
WRITE filename,RZERO
is used in a program, this operand is required.

BLKSIZE=n: This operand indicates the size of the I/O area by specifying the maximum number of

characters that are transferred to or from the area at any one time. When undefined, variable length or spanned records are read or written, the area must be large enough to accommodate the largest record.

For details on how to compute n, see *DOS/VSE Macro User's Guide*, as listed in the Preface.

IOCS uses this specification to construct the count field of the CCW for reading or writing records.

CONTROL=YES: Include this operand if a CNTRL macro is issued for this file. The CNTRL macro for seeking on a disk allows you to specify a track address on which access movement should begin for the next READ or WRITE macro. While the arm is moving, you may process data and/or request I/O operations on other devices.

DEVADDR=SYSnnn: This operand must specify the symbolic unit (SYSnnn) associated with a file if the symbolic unit is not provided via an EXTENT job control statement. If such a unit is provided, its specification overrides the DEVADDR parameter. This specification, or symbolic unit, represents an actual I/O address and is used in the ASSGN job control statement to assign the actual I/O device address to the file.

Note: EXTENT job control statements provided for DAM must be supplied in ascending order, and the symbolic units for multivolume files must be assigned in consecutive order.

DEVICE= {2311|2314|3330|3340|3350}: This operand specifies the device on which the file is located. If this operand is omitted, 2311 is assumed.

For devices supported by DOS/VSE and not included in the above operand specification, specify device codes as listed in Figure 3-6.

Note: DAM does not permit the use of different size data modules (on a 3340) for a multivolume file.

DSKXTNT=n: This operand indicates the maximum number of extents (up to 256) that are specified for a file. When this operand is used together with FIXUNB, VARUNB, or UNDEF specified in the RECFORM operand, it indicates that a relative ID is used in the SEEKADR and IDLOC locations. If DSKXTNT=n is omitted, a physical ID is assumed in the SEEKADR and IDLOC locations.

If RECFORM=SPNUNB is specified, DSKXTNT is required. If relative addressing is used, the RELTYPE operand must also be specified.

ERRBYTE=name: This operand is required for IOCS to supply indications of exceptional conditions to your program. The name of a 2-byte field (in which IOCS can store the error-condition or status codes) is entered.

ERREXT=YES: This operand enables unrecoverable I/O errors (occurring before a data transfer takes place) to be indicated to your program. This error information is indicated in the bytes named in the ERRBYTE operand and is available after the WAITF macro has been issued.

FEOVD=YES: This operand is specified if code is generated to handle end-of-volume records. It should be specified only when reading a file which was built using DTFSD and the FEOVD macro.

HOLD=YES: This operand can be specified only if the track hold function is

- specified at system generation time,
- included in the DAMOD macro, and
- used when the file is referenced.

IDLOC=name: This operand is included if you want IOCS to supply the ID of a record after each READ or WRITE (ID or KEY) is completed. Specify the name of a record reference field in which IOCS is to store the ID. WAITF should be used before referencing this field. Do not specify the same field for IDLOC and SEEKADR.

DEVICE = specification	Device in use							
	2311	2314	2319	3330-1,2*	3330-11**	3340, 35MB	3340, 70MB	3350
Default	x				x			x
2311	x				x			x
2314		x	x		x			x
3330				x	x			x
3340					x	x	x	x
3350					x			x

* Also 3350 in 3330-1 compatibility mode.
** Also 3350 in 3330-11 compatibility mode.

Figure 3-6. DEVICE= specifications for DTFDA.

IOAREA1=name: This operand must be included to specify the name of the input/output area used for the file. IOCS routines transfer records to or from this area. The specified name must be the same as the name used in the DS instruction that reserves this area of storage.

KEYARG=name: This operand must be included if records are identified by key; that is, if either of the macros

    READ filename,KEY
    WRITE filename,KEY

is used in a program, this entry and the corresponding KEYLEN operand are required. KEYARG specifies the name of the key field in which you supply the record key to IOCS.

The KEYARG operand is required for formatting WRITE (WRITE filename,AFTER) operations for files containing keys if RECFORM=VARUNB or SPNUNB. It is required also when the macro

    READ filename,ID

is specified and if KEYLEN is not zero. When record reference is by key, IOCS uses this specification at assembly time to construct the data address field of the CCW for search commands.

KEYLEN=n: This operand must be included if record reference is by key or if keys are read or written. It specifies the number of bytes in each key. All keys must be the same length. If this operand is omitted, IOCS assumes a key length of zero.

If there are keys recorded on DASD and this entry is absent, a WRITE ID or READ ID writes or reads the data portion of the record.

When record reference is by key, IOCS uses this specification to construct the count field of the CCW for this file. IOCS also uses this in conjunction with IOAREA1 to determine where the data field in the I/O area is located.

LABADDR=name: You may require one or more user labels in addition to the standard file label. If so, you must include your own routine to check, or write, the labels. The name of such a routine is specified in this operand. IOCS branches to this routine after it has processed the standard label.

MODNAME=name: This operand specifies the name of the logic module that is used with the DTF table to process the file. If the logic module is assembled with the program, MODNAME must specify the same name as the DAMOD macro. If this entry is omitted, standard names are generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each task should have its own uniquely defined save area. Each time an imperative macro (except OPEN, OPENR, or LBRET) is issued, register 13 must contain the address of the save area associated with the task. The fact that the save areas are unique for each task makes the module reentrant, that is, capable of being used concurrently by several tasks.

READID=YES: This operand must be included if, in your program, the macro READ filename, ID is used.

READKEY=YES: This operand must be included if, in your program, the macro READ filename,KEY is used.

RECFORM= {FIXUNB|SPNUNB|UNDEF|
VARUNB}:
This operand specifies the type of records in the input or output file. The specifications are:

FIXUNB
> For fixed-length records. All records are considered unblocked. If you want blocked records, you must provide your own blocking and deblocking.

SPNUNB
> For spanned records. This specification is for unblocked variable-length logical records of less than 32,768 bytes per record.

UNDEF
> For undefined records. This specification is required only if the records are of undefined format.

VARUNB
> For variable-length records. This specification is for unblocked variable-length records.

For a definition of record formats see *DOS/VSE Data Management Concepts*, as listed in the Preface.

RECSIZE=(r): This operand must be included if undefined records are specified (RECFORM=UNDEF). It specifies the number of the general-purpose register (any of 2 through 12) that contains the length of each individual input or output record.

Whenever an undefined record is read, IOCS supplies the length of the data area for that record in the specified register.

When an undefined record is written, you must load the length of the data area of the record (in bytes) into this register, before you issue the WRITE

macro for the record. IOCS adds the length of the key when required.

When records are written (AFTER specified in the WRITE macro), IOCS uses the length to construct the count area written on DASD. IOCS adds the length of both the count and the key when required.

RELTYPE= {DEC|HEX}: This operand specifies whether the zoned decimal (DEC) or hexadecimal (HEX) form of the relative ID is to be used. When FIXUNB, VARUNB, or UNDEF is specified in the RECFORM operand, RELTYPE should be supplied only if the DSKXTNT operand (relative ID) is specified. If omitted, a hexadecimal relative ID is assumed. However, if DSKXTNT is also omitted, a physical ID is assumed in the SEEKADR and IDLOC addresses.

If RECFORM=SPNUNB is specified, the RELTYPE operand is required when relative addressing is used. If RELTYPE is omitted, a physical ID is assumed in the SEEKADR and IDLOC addresses.

SEEKADR=name: This operand must be included to specify the name of your track-reference field. In this field, you store the track location of the particular record read or written. IOCS refers to this field to determine which volume and which track contains the desired record. Whenever records are to be located by searching for a specified ID, the track-reference field must also contain the number of the record on the track.

SEPASMB=YES: Include this operand only if the DTFDA will be assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

SRCHM=YES: If records are identified by key, this operand may be included to cause IOCS to search multiple tracks for each specified record. The macros
    READ filename,KEY
    WRITE filename,KEY
cause IOCS to search the track specified in the track-reference field and all following tracks in the cylinder, until the record is found or the end of the cylinder is reached. If the file ends before the end of the cylinder and the record is not found, the search continues into the next file, if any, on the cylinder. EOC, instead of NRF, is indicated. Without SRCHM=YES, each search is confined to the specified track.

TRLBL=YES: This operand, if specified with the LABADDR operand, indicates that user standard trailer labels are to be read or written following the user standard header labels on the user label track. Both operands must be specified for trailer label processing. For more information on processing labels, see "Label Processing" in *DOS/VSE Macro User's Guide*, as listed in the Preface.

TYPEFLE= {INPUT|OUTPUT}: This operand must be included to indicate how standard volume and file labels are to be processed. INPUT indicates that standard labels are to be read; OUTPUT indicates that standard labels are to be written.

This entry is always required.

VERIFY=YES This operand is included if you want to check the parity of disk records after they are written. If this operand is omitted, any records written on a disk are not verified.

WRITEID=YES: This operand must be included if the DASD storage location for writing any output record or updating an input file is specified by a record ID (identifier); that is, whenever the macro
    WRITE filename,ID
is used in the program, this operand is required.

WRITEKY=YES: This operand must be included if the DASD location for writing any output record or updating an input file is specified by record key, that is, whenever
    WRITE filename,KEY
is used.

XTNTXIT=name: This operand is included if you want to process label extent information. It specifies the name of your extent exit routine. During an OPEN, IOCS branches to your routine after each specified extent is checked. Upon entering your routine, IOCS stores, in register 1, the address of a 14-byte field that contains the label extent information (in binary form) retrieved from the label information cylinder. If user labels are present, the user label track is returned as a separate extent and the lower limit of the first normal extent is increased by one track. The format of this field is shown in Figure 3-7. Return to IOCS by use of the LBRET macro. Registers 2 through 13 are available in the XTNTXIT routine. Within the routine you cannot issue a macro that calls a transient routine (such as OPEN, OPENR, CLOSE, CLOSER, DUMP, PDUMP, CANCEL, CHKPT, etc.).

Bytes	Contents
0	Extent type code (as specified in the extent statement)
1	Number of extent (as determined by the extent card sequence)
2-5	Lower limit of the extent (cchh)
6-9	Upper limit of the extent (cchh)
10-11	Symbolic unit number (in hexadecimal format)
12-13	Not used

Figure 3-7 Label extent information field

DAMOD or DAMODV Macro

Operation	Operand	Remarks
DAMOD[1] or DAMODV[2] Must be included	AFTER=YES	When WRITE with the operand AFTER or RZERO is used
	ERREXT=YES	Required if non-data-transfer error conditions are to be indicated in the ERRBYTE status bits
	FEOVD=YES	Required if support for sequential disk end-of-volume records is desired
	HOLD=YES	Required if the track hold function is to be used
	IDLOC=YES	Required if IDLOC specified in DTFDA
	RDONLY=YES	Required if a read only module is to be generated
	RECFORM= {FIXUNB[1]| UNDEF[1]| VARUNB[2]| SPNUNB[2]}	Describes record format
	RELTRK=YES	Required if the module is to process relative identifiers along with physical identifiers
	RPS=SVA	To assemble RPS logic modules
	SEPASMB=YES	If the module is assembled separately

[1]. DAMOD is for fixed length unblocked and undefined records.
[2]. DAMODV is for variable length and spanned unblocked records.

Figure 3-8. DAMOD macro.

AFTER=YES: This operand generates a logic module that can perform a formating WRITE (count, key, and data). It performs the functions required by WRITE filename,AFTER; and WRITE filename,RZERO. The module also processes any files in which the AFTER operand is not specified in the DTF.

HOLD=YES: This operand is specified if the track hold function is

- specified at system generation time, and
- included in the DTFDA macro, and
- used when the file is referenced.

For more information see the DTFDA HOLD operand.

ERREXT=YES: Include this operand if unrecoverable I/O errors (occurring before a data transfer takes place) are to be indicated to your program in the bytes named in the DTF ERRBYTE operand.

FEOVD=YES: This operand is specified if coding is to handle end-of-volume records. It should be specified only if you are reading a file built using DTFSD and the FEOVD macro.

IDLOC=YES: This operand generates a logic module that returns record identifier (ID) information to you. The module also processes any files in which the IDLOC operand is not specified in the DTF.

RDONLY=YES: This operand causes a read-only

module to be generated. Whenever this operand is specified, any DTF used with the module must have the same operand.

RECFORM= {FIXUNB\|SPNUNB\|UNDEF\| VARUNB}:
If UNDEF is specified, the logic module generated can handle both unblocked fixed-length and undefined records. If the operand is omitted or if FIXUNB is specified, the logic module generated can handle only fixed-length unblocked records. If SPNUNB is specified, the module can handle both format V (variable length) and spanned format records. If VARUNB is specified, the module can handle only format V records.

RELTRK=YES: This operand generates a logic module that can process with both physical and relative identifiers. If the operand is omitted, the module can process only with physical identifiers.

RPS=SVA: This operand causes the RPS logic modules to be assembled.

SEPASMB=YES: Include this operand only if the module will be assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and the module name to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the problem program and no CATALR card is punched.

Standard DAMOD Names

Each name begins with a 3-character prefix (IJI) and continues with a 5-character field corresponding to the options permitted in the generation of the module.

DAMOD name = IJIabcde

```
a  = F   RECFORM=FIXUNB
   = B   RECFORM=UNDEF (handles both UNDEF and
           FIXUNB)
   = S   RECFORM=SPNUNB
   = V   RECFORM=VARUNB

b  = A   AFTER=YES,RPS=SVA omitted
   = W   AFTER=YES, RPS=SVA
   = Z   neither is specified

c  = E   IDLOC=YES and FEOVD=YES
   = I   IDLOC=YES
   = R   FEOVD=YES
   = Z   neither is specified

d  = H   ERREXT=YES and RELTRK=YES
   = P   ERREXT=YES
   = R   RELTRK=YES
   = Z   neither is specified

e  = W   HOLD=YES and RDONLY=YES
   = X   HOLD=YES
   = Y   RDONLY=YES
   = Z   neither is specified
```

Subset/Superset DAMOD Names

Figure 3-9 shows the subsetting and supersetting allowed for DAMOD names. Five parameters allow supersetting. For example, the module IJIBAIZZ is a superset of the module with the name IJIFAZZZ.

```
          +  +  +  +  +
I  J  I   B  A  E  H  X
          F  Z  I  P  Z
          +     Z  Z  +
          S     +  +  W
          V     E  H  Y
                R  R
                Z  Z
```

+ Subsetting/supersetting permitted.

Figure 3-9. Subsetting and supersetting of DAMOD names.

Notes:

1. The module IJIBAEHW will cause assembly error message IPK154 TOO MANY ENTRY SYMBOLS
The valid entry points for this module total more than 100, which is the maximum for assembler language. Specify less parameters for DAMOD if you can. Otherwise, you must assemble your own module for your program.
2. Your program can have only one DAMOD for fixed unblocked or undefined records and/or only one DAMOD for variable unblocked or spanned unblocked records; otherwise, duplicate name flagging occurs during assembly time.

DTFDI Macro
The DTFDI macro provides device independence for system logical units.

M	DEVADDR=SYSxxx	(SYSIPT, SYSLST, SYSPCH, or SYSRDR). System logical unit.
M	IOAREA1=xxxxxxxx	Name of first I/O area.
O	CISIZE=n	Size of FBA DASD control interval.
O	EOFADDR=xxxxxxxx	Name of your end-of-file routine.
O	ERROPT=xxxxxxxx	(IGNORE, SKIP, or name of your error routine). Prevents termination on errors.
O	FBA=YES	Specifies a Fixed Block Architecture device.
O	IOAREA2=xxxxxxxx	If two I/O areas are used, name of second area.
O	IOREG=(r)	Register number. If omitted and 2 I/O areas are used, register 2 is assumed. General registers 2-12, written in parentheses.
O	MODNAME=xxxxxxxx	DIMOD name for this DTF. If omitted, IOCS generates a standard name. Ignored with FBA DASD or 3800 advanced printer buffering.
O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module.
O	RECSIZE=nnn	Number of characters in record. Assumed values: 121 (SYSLST), 81 (SYSPCH), 80 (otherwise).
O	SEPASMB=YES	DTFDI to be assembled separately.
O	TRC=YES	For 3800, output data lines include table reference character.
O	WLRERR=xxxxxxxx	Name of your wrong length record routine.

M=Mandatory
O=Optional

Figure 3-10. DTFDI macro operands.

DEVADDR= {SYSIPT|SYSLST|SYSPCH| SYSRDR}:
This operand must specify the symbolic unit associated with this system file. Only the system names shown above may be specified. The logical device SYSLST must not be assigned to the 2560 or 5424/5425.

IOAREA1=name: This operand must specify the name of the input or output area used with the file. The input and/or output routines transfer records to or from this area.

If the DTFDI macro is used to define a printer file, or a card file to be processed on a 2540, 2560, 3525, or 5424/5425, the first byte of the output area must contain a control character.

CISIZE=n: This operand specifies the FBA control interval size. The value n must be an integral multiple of the FBA logical block size and, if greater than 8K, must be a multiple of 2K. The maximum value is 32768 (32K), except when assigned to SYSLST or SYSPCH, when the maximum is 30720 (30K).

If FBA=YES is specified, and CISIZE is omitted, CISIZE=0 is assumed. Control interval size may be overridden for an output file at execution time by specifying the CISIZE parameter of the DLBL control statement. For an input file, the CISIZE value in the format-1 label is used.

EOFADDR=name: This operand specifies the name of your end-of-file routine. It is required only if SYSIPT or SYSRDR is specified.

IOCS branches to this routine when it detects an end-of-file condition. In this routine, you can perform any operations necessary for the end-of-file condition (you generally issue the CLOSE or CLOSER macro).

IOCS detects the end-of-file condition by recognizing the characters /* in positions 1 and 2 of the record for cards, a tapemark for tape, and end-of-file record for disk. If the system logical units SYSIPT and SYSRDR are assigned to a 5424/5425, IOCS requires that the /* card, indicating end-of-file, be followed by a blank card. An error condition results if the records are allowed to run out without a /* card (and without a /& card, if end-of-job). IOCS detects the end-of-file condition on diskette units by recognizing that end-of-data has been reached on the current volume and that there are no more volumes available.

ERROPT= {IGNORE|SKIP|name}: This operand does not apply to output files. For output files for most devices, the job is automatically terminated after IOCS has attempted to retry writing the record; for 2560 or 5424/5425 output files, normal error recovery procedures are followed.

This operand applies to wrong-length records if WLRERR is omitted. If both ERROPT and WLRERR are omitted and wrong-length records occur, IOCS ignores the error.

ERROPT specifies the function to be performed for an error block. If an error is detected when reading a magnetic tape, or a disk or a diskette volume, IOCS attempts to recover from the error. If the error is not corrected, the job is terminated unless this operand is included to specify other procedures to be taken. The three specifications are described below.

IGNORE
> indicates that the error condition is to be ignored. The address of the error record is made available to you for processing (see *CCB Macro*).

SKIP
> indicates that the error block is not to be made available for processing. The next record is read and processing continues.

name
> indicates that IOCS is to branch to your routine when an error occurs, where you may perform whatever functions are desired or simply note the error condition. The address of the error record is supplied in register 1. The contents of the IOREG register may vary and should not be used for error records. Also, you must not issue any GET instructions in your error routine. If you use any other IOCS macros, you must save the contents of register 14. If RDONLY=YES is specified, or if the file is assigned to an FBA device, you must also save the contents of register 13. At the end of the error routine, return to IOCS by branching to the address in register 14. The next record is then made available for processing.

FBA=YES: Specifies a Fixed Block Architecture device. If used, MODNAME specification is ignored and an IBM-supplied module is used.

IOAREA2=name: Two input or output areas can be allotted for a file to permit overlapped GET or PUT processing. If this operand is included, it specifies the name of the second I/O area.

IOREG=(r): When two I/O areas are used, this operand specifies the general purpose register (any of 2 through 12) that points to the address of the next record. For input files, it points to the logical record available for processing. For output files, it points to the address of the area where you can build a record. If omitted, and two I/O areas are used, register 2 is assumed.

MODNAME=name: This operand may be used to specify the name of the logic module used with the DTF table to process the file. If the logic module (DIMOD) is assembled with the program, the MODNAME parameter in this DTF must specify the same name as the DIMOD macro.

If this entry is omitted, standard names are generated for calling the logic module. If two different DTF macros call for different functions that can be handled by a single module, only one standard-named module is called.

RDONLY=YES: This operand is specified if the DTF is to be used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each task should have its own uniquely defined save area, and each time an imperative macro (except OPEN, OPENR or LBRET) is issued, register 13 must contain the address of the save area associated with that task. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If an ERROPT or WLRERR routine issues I/O macros using the same read-only module that caused control to pass to either error routine, the program must provide another save area. One save area is used for the initial I/O operations, and the second for I/O operations in the ERROPT or WLRERR routine. Before returning to the module that entered the error routine, register 13 must be set to the save area address originally specified for the task.

If the operand is omitted, the module generated is not reenterable and no save area need be established.

RECSIZE=n: This operand specifies the length of the record. For input files (SYSIPT and SYSRDR), the maximum allowable record size is 80 bytes. For output files, RECSIZE must include one byte for control characters. The maximum length specification is 121 for SYSLST and 81 for SYSPCH.

For disk files, 121 must be specified for SYSLST, and 81 for SYSPCH to assure device independence. For printers and punches, DIMOD assumes a S/370-type control character if the character is not a valid ASA character. The program checks ASA control characters before S/370-type control characters. Therefore, if it is a valid ASA control character (even though it may also be a S/370-type control character), it is used as an ASA control character. Otherwise, it is used as a S/370-type control character.

Control character codes are listed in *DOS/VSE Macro User's Guide*; note, however:

- 2520 stacker selection codes must be used for the 1442.
- 2540 stacker selection 3 must not be used if device independence is to be maintained.

If this operand is omitted, the following is assumed:

80 bytes for SYSIPT.
80 bytes for SYSRDR.
81 bytes for SYSPCH.
121 bytes for SYSLST.

The use of assumed values for the RECSIZE operand assures device independence. For disk and diskette files, the assumed values are required to assure device independence.

SEPASMB=YES: Include this operand only if the DTFDI will be assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and defines the filename as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TRC=YES: This operand applies to the IBM 3800 Printing Subsystem. TRC=YES specifies that a table reference character is included as the first byte of each output data line (following the optional print control character). The printer uses the table reference character (0, 1, 2, or 3) to select the character arrangement table corresponding to the order in which the table names have been specified with the CHARS parameter on the SETPRT job control statement (or SETDRT macro instruction).

If the device allocated is not a printer and TRC=YES is specified, the table reference character is treated as data when a PUT is issued. If the device is a non-3800 printer, the table reference character is removed and not printed.

WLRERR=name: This operand applies only to input files on devices other than diskette units. It specifies the name of your routine to which IOCS branches if a wrong-length record is read on a tape or disk device.

DIMOD Macro

Listed here are the operands you can specify for
DIMOD. The header card contains DIMOD in the
operation field and may contain a module name in
the name field. If the module is omitted, IOCS gener-
ates a standard module name.

IOAREA2=YES: Include this operand if a second
I/O area is needed. A module with this operand can
be used with DTFDIS specifying either one or two I/O
areas. If the operand is omitted or is invalid, one I/O
area is assumed.

RDONLY=YES: This operand causes a read only
module to be generated. Whenever this operand is
specifed, any DTF used with the module must have
the same operand.

RPS=SVA: This operand causes the RPS logic mo-
dules to be assembled.

SEPASMB=YES: Include this operand only if the
module is assembled separately. This causes a
CATALR card with the module name (standard or
user-specified) to be punched ahead of the object
deck and defined the module name as an ENTRY
point in the assembly. If the operand is omitted, the
assembler assumes that the module is being assem-
bled with the DTF and the problem program and no
CATALR card is punched.

TRC=YES: Include this operand to specify wheth-
er the module is to test the table reference character
indicator in the DTFDI or ignore that indicator. If
TRC=YES is specified, the generated module can
process output files with table reference characters
and those without. If the TRC operand is specified,
TYPEFLE=INPUT must not be specified.

TYPEFLE= {OUTPUT|INPUT}: Include this
operand to specify whether the module is to process
input or output files. If OUTPUT is specified, the
generated module can process both input and output
files.

Standard DIMOD Names

Each name begins with a 3-character prefix (IJJ)
followed by a 5-character field corresponding to the
options permitted in the generation of the module.

DIMOD name = IJJabcde

a	= F	
b	= C	RPS=SVA is not specified
	= V	RPS=SVA
c	= B	TYPEFLE=OUTPUT (both input and output)
	= I	TYPEFLE=INPUT
d	= I	IOAREA2=YES
	= Z	IOAREA2=YES is not specified
e	= C	RDONLY=YES
	= D	RDONLY=YES is not specified

Subset/Superset DIMOD Names

Figure 3-11 illustrates the subsetting and superset-
ting allowed for DIMOD names. All of the operands
except TRC=YES allow subsetting. A module name
specifying B is a superset of the module specifying I:
for example. IJJFFCBID is a superset of the module
IJJFCIID.

The IBM-supplied preassembled logic modules do
not have TRC=YES. The system programmer can
reassemble them with TRC=YES to support 3800
table reference characters. Although the code that is
generated for a module assembled with TRC=YES is
different from the code that is generated for a mo-
dule with TRC=NO, the module name is the same. If
some, but not all DIMOD logic modules are reassem-
bled this way, it may interfere with subsetting or
supersetting.

```
                 +  +  *
      I  J  J  F  C  B  I  C
                     I  Z  D
```
+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.
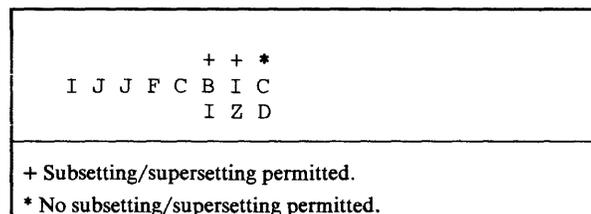
Figure 3-11. Subsetting and supersetting of DIMOD names.

DTFDR Macro

You must use the DTFDR macro to define each 3886 file in your program.

M	DEVADDR=SYSxxx	Symbolic unit assigned to 3886 optical character reader.
M	FRNAME=xxxxxxxx	Phase name of format record to be loaded upon file opening.
M	FRSIZE=nn	Number of bytes to be reserved in DTF expansion for format records.
M	EXITIND=xxxxxxxx	Name of completion code return area.
M	IOAREA1=xxxxxxxx	Name of file input area.
M	HEADER=xxxxxxxx	Name of area for header record from 3886.
M	EOFADDR=xxxxxxxx	Address of your end-of-file routine.
M	COREXIT=xxxxxxxx	Name of your error condition routine.
O	DEVICE=3886	If omitted, 3886 is assumed.
O	RDONLY=YES	If DTF is to be used with read-only module.
O	MODNAME=xxxxxx	Name of DRMODxx logic module for this DTF. If omitted, IOCS generates standard name.
O	BLKSIZE=nnn	Length of area named by IOAREA1. If omitted, the maximum length of 130 is assumed.
O	SEPASMB =YES	If DTFDR is to be assembled separately.
O	SETDEV=YES	If SETDEV macro is issued in your program to load a different format record into the 3886.

M=Mandatory
O=Otional

Figure 3-12. DTFDR macro operands.

DEVADDR=SYSnnn: Specifies the symbolic unit to be associated with the logical file. The symbolic unit (SYSnnn) is associated with an actual I/O device through the job control ASSGN statement.

FRNAME=phasename: Specifies the phase name of the format record to be loaded when the file is opened.

FRSIZE=number: Specifies the number of bytes to be reserved in the DTF expansion for format records. The number must equal at least the size of the largest DFR macro expansion and its associated DLINT macro expansions, plus four. This size is printed in the ninth and tenth bytes of the DFR macro expansion.

If you use the SETDEV macro in your program to change format records, you can reduce the library retrieval time by specifying a size large enough to contain all the frequently used format records. The area should then be equal to the sum of the format record sizes, plus four bytes for each format record. When the SETDEV macro is issued, the format record is loaded into this area from the core image library if it is not already present in the area.

EXITIND=name: Specifies the symbolic name of the 1-byte area in which the completion code is returned to the COREXIT routine for error handling from an I/O operation.
The completion codes are:

Code		Meaning
Dec	**Hex**	
240	X'F0'	No errors occurred. (This code should not be present when the COREXIT routine receives control.)
241	X'F1'	Line mark station timing mark check error.
242	X'F2'	Nonrecovery error. Do not issue the CNTRL macro to eject the document from the machine. Have the operator remove the document.
243	X'F3'	Incomplete scan.
244	X'F4'	Line mark station timing mark check and equipment check.
249	X'F9'	Permanent error.

Note: If any of these errors occur while the file is being opened, the COREXIT routine does not receive control and the job is canceled.

IOAREA1=name: Specifies the symbolic name of the input area to be used for the file. The area must be as large as the size specified in the BLKSIZE parameter. If BLKSIZE is not specified, the input area must be 130 bytes.

HEADER=name: Specifies the symbolic name of the 20-byte area to receive the header record from the 3886.

EOFADDR=name: Specifies the symbolic address of your end-of-file routine. LIOCS branches to this routine whenever end of file is detected on the 3886.

COREXIT=name: Provides the symbolic name of your error correction routine. LIOCS branches to this routine whenever an error is indicated in the EXITIND byte.

You can attempt to recover from various errors that occur on the 3886 through the COREXIT routine you provide. Your COREXIT routine receives control whenever one of the following conditions occurs:

- Incomplete scan
- Line mark station timing mark check error
- Nonrecovery error
- Permanent error

Note: If any of these errors occur while the file is being opened, the COREXIT routine does not receive control and the job is canceled.

Figure 3-13 describes normal functions for the COREXIT routine for the various error conditions and provides the exits that must be taken from the COREXIT routine.

Error messages are provided to describe errors to the operator during program execution.

DEVICE=3886: Indicates that 3886 is the I/O device for this file. This operand may be omitted.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each DTF should have its own uniquely defined save area.

Each time an imperative macro (except OPEN or OPENR) is issued using a particular DTF, register 13 must contain the address of the save area associated with that DTF. The fact that the save areas are unique or different for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If a COREXIT routine issues I/O macros using the same read-only module that caused control to pass to either error routine, your program must provide another save area. One save area is used for the normal I/O operations, and the second for I/O operations in the COREXIT routine. Before returning to the module that entered the COREXIT routine, register 13 must contain the save area address originally specified for that DTF.

If this operand is omitted, the module generated is not reenterable, and no save area is required.

MODNAME=name: This operand may be used to specify the name of the logic module used with the DTF table to process the file. If the logic module (DRMOD) is assembled with the program, the MODNAME parameter in this DTF must specify the same name as the DRMOD macro.

If this entry is omitted, standard names are generated for calling the logic module. If two different DTF macros call for different functions that can be handled by a single module, only one standard-named module is called.

BLKSIZE=nnn: Specifies the length of the area named by the IOAREA1 keyword. The length of the area must be equal to the length of the longest record to be passed from the 3886.

If this operand is omitted, the maximum length of 130 is assumed.

Note: DOS/VSE LIOCS does not allow you to block records read from the 3886.

SEPASMB=YES: Specifies that the DTF will be assembled separately. If this operand is specified, a CATALR card with the filename is punched before

Error	Normal COREXIT Function	Exit to
X'F2'	Eliminate the data that has been read from this document and prepare to read the next input document (See Note 1).	Routine in your program to read the next document.
X'F4' or X'F9'	Do whatever processing is necessary before the job is canceled. (See Note 1).	Your end-of-job routine.
X'F1'	Do any processing that may be required. The document may have been read incorrectly; you may want to delete all data records from the document (see Note 2).	Branch to the address in register 14 to return to the instruction following the macro causing the error.
X'F3'	Rescan the line using another format record or using image processing and editing the record in your program (see Note 2).	Branch to the address in register 14 to return to the instruction following the macro causing the error.
Note 1: If in your COREXIT routine, you issue an I/O macro to the 3886 and an error occurs during that operation, control is returned to the beginning of the COREXIT routine. You must take precautions in the COREXIT routine to prevent looping in this situation. If no errors occur, control returns to the instruction following the I/O macro. **Note 2:** If, in your COREXIT routine, you issue an I/O macro to the 3886, control always returns to the instruction following the macro. You should then check the completion code to determine the outcome of the operation.		

Figure 3-13. COREXIT routine functions.

the deck and the filename is defined as an ENTRY point for the assembly.

SETDEV=YES: Specifies that the SETDEV macro is issued in your program to load a different format record into the 3886.

DRMOD Macro

Listed here are the operands you can specify for DRMOD. The first card contains DRMOD in the operation field and may contain a module name in the name field.

DEVICE=3886: Specifies that the 3886 is the input device. This operand may be omitted.

SEPASMB=YES: Must be specified if the I/O module is assembled separately. This entry causes a CATALR card to be punched preceding the module.

RDONLY=YES: This operand generates a read only module. RDONLY=YES must also be specified in the DTF. For additional programming requirements concerning this operand, see the DTFDR RDONLY operand.

SETDEV=YES: Is specified if the SETDEV macro may be used when processing a file with this I/O module. If SETDEV=YES is specified in the DRMOD macro but not in the DTFDR macro, the SETDEV macro cannot be used when processing that file.

Standard DRMOD Names

Each name consists of eight characters. They are: IJMZxxD0. The fifth and sixth characters are variables as follows:

- If SETDEV=YES is specified, the fifth character is S; otherwise it is Z.

- If RDONLY=YES is specified, the sixth character is R; otherwise it is Z.

Note: Subsetting/supersetting is allowed with the SETDEV keyword, but not with the RDONLY keyword.

DFR Macro

The DFR macro defines attributes common to a group of line types.

M	FONT=xxxx	Default font for all codes described by format record.
O	REJECT=x	Replacement character for any reject character in the data record read by the 3886. If omitted, X'3F' is assumed.
O	ERASE=YES	Group and character erase symbols are to be recognized. If omitted, NO is assumed.
O	CHRSET=n	Specifies recognizing character (see Figure 3-15). If omitted, 0 is assumed.
O	EDCHAR=(x, ..)	Characters that may be deleted from any field that is read. If omitted, no character deletion occurs.
O	BCH=n	Batch numbering is to be preformed by 3886. If specified, BCHSER is invalid.
O	BCHSER=n	Both batch and serial numbering are to be performed. If specified, BCH is invalid.
O	NATNHP=YES	European Numeric Hand Printing (ENHP) characters 1 and 7 are used. If omitted, NO is assumed, indicating that Numeric Hand Printing (NHP) character 1 + 7 are used.

M=Mandatory
O=Optional

Figure 3-14. DFR macro operands.

FONT =code: Specifies the default font for all fields described by the format record. The default font is used to read a field unless another font is specified for an individual field through the DLINT macro. This is the only required operand in the DFR macro. The valid codes and the fonts they represent are:

NUMA	Numeric OCR-A font
ANA1	Alphameric OCR-A font (mode 1)
ANA2	Alphameric OCR-A font (mode 2)
NUMB	Numeric OCR-B font (mode 3)
ANB1	Alphameric OCR-B font
NHP1	Numeric hand printing (normal mode)
NHP2	Numeric hand printing (verify mode)
GOTH	Gothic font
MRKA	Mark OCR-A font
MRKB	Mark OCR-B font

For a description of these fonts, see the appropriate IBM 3886 device manuals.

REJECT = character: Indicates the character that is to be substituted in the data record for any reject character read by the device. If this operand is omit-ted, X'3F' is assumed. Reject characters are charac-ters that are not recognizable by the device.

Note: This note applies to the keywords REJECT and ED-CHAR. Aphostrophes enclosing the character are optional for all characters except special characters used in macro operands. For a description of these characters, see *OS/VS-DOS/VSE-VM/370 Assembler Language*, as listed in the Preface.

ERASE = {YES|NO}: Specifies whether group and character erase symbols are to be recognized as valid symbols. If this operand is not specified, NO is assumed. For more information on group and char-acter erase symbols, see the appropriate IBM 3886 device manuals.

CHRSET = {0|1|2|3|4|5}: Specifies which one of the options in Figure 3-15 is to be used for recogniz-ing characters. If this operand is not entered, 0 is assumed.

OCR-A			OCR-B			
Numeric Mode	Alphameric Modes		Numeric Mode	Alphameric Mode		
Highspeed Printers or Typewriters	Mode 1 (Highspeed) Printer)	Mode 2 (Typewriter)	Highspeed Printers or Typewriters		Hexa-decimal Code	Format Record Codes
≢	≢	≢	$	$	5B	00
£	£	£	£	£	5B	01
¥	¥	¥	¥	¥	5B	02
	Ñ	Ñ		Ñ	7B	
≢	≢	≢	$	$	5B	03
	Å	Å		Å	5B	
	Æ	Æ		Æ	7B	
	Ø	Ø		Ø	7C	04
≢	≢	≢		Ü Note	5B	
	Ä	Ä		Ä	7B	
	Ö	Ö		Ö	7C	
	Ü	Ü			F0	05

Note: In OCR-A font the Ü is coded as a zero and should be used only in alphabetic fields.

Figure 3-15. Character set option list.

EDCHAR = (x,...): Specifies up to six characters that may be deleted from any field that is read. The EDCHAR parameter in the EDITn keyword of the DLINT macro controls this function for individual fields. If this operand is omitted, no character deletion is performed. See the note under the REJECT operand discussion for characters that must be specified in quotes. For example, to specify the characters &, >, and), you would code EDCHAR=('&','>',')').

BCH = {1|2|3}: Indicates that batch numbering is to performed by the 3886. Specifying 1, 2, or 3 indicates that documents routed to a stacker are to be batch numbered. Specifying 1 indicates stacker A, 2 indicates stacker B, 3 indicates both stackers. If this operand is specified, the BCHSER operand is invalid. If neither BCH nor BCHSER are entered, no batch numbering is performed. This operand is valid only if the serial numbering feature is installed on the 3886. For more information on batch numbering, see the appropriate IBM 3886 device manuals.

RCHSER = {1|2|3}: Indicates that both batch and serial numbering are to be performed by the 3886. Specifying 1, 2, or 3 indicates that documents routed to a stacker are to be batch and serial numbered. Specifying 1 indicates stacker A, 2 indicates stacker B, 3 indicates both stackers. If this operand is specified, the BCH operand is invalid. If neither BCH nor BCHSER is specified, batch and serial numbering are not performed. This operand is valid only if the serial numbering feature is installed on the 3886. For more information on batch and serial numbering, see the appropriate IBM 3886 device manuals.

NATNHP = {YES|NO}: Specifies which of the numeric hand printing character set options are used for the numbers 1 and 7. YES indicates that the European Numeric Hand Printing (ENHP) characters 1 and 7 are used; NO indicates the Numeric Hand Printing (NHP) characters 1 and 7 are used. If this operand is not entered, NO is assumed.

DLINT Macro

The DLINT macro describes one line type in a format group and the individual fields in the line.

M	LFR=nn	Line format record for the line.
M	LINBEG=nn	Specifies beginning of a line.
O	IMAGE=YES	Data record is to be in image mode. If omitted, NO (standard mode) is assumed.
O	NOSCAN=(n,n)	Indicates an area on the document line that is to be ignored by the 3886.
O	FLDn=(n,n,NCRIT,xxx)	Describes a field in a line. n in the FLD keyword may be from 1 to 14; if specified, a corresponding EDITn keyword must follow each FLDn keyword.
O	EDITn=(xxxxxx,EDCHAR)	Specifies editing functions to be performed on the data by 3886. A corresponding FLDn keyword must precede each EDITn keyword.
O	FREND=YES	Indicates last DLINT macro for the format record. If omitted, NO is assumed meaning that further DLINT macros follow.

M=Mandatory
O=Optional

Figure 3-16. DLINT macro operands.

Line Information Entries

LFR=number This operand is required. It specifies the line format record number for the line. The decimal number specified must be in the range of 0 through 63.

The line format record describes the format of one type of line; the line format record number is used to identify the line format record. This number is specified in the READ macro when you read a line of data from a document.

LINBEG=number: This operand is required. It specifies the beginning of a line. The beginning position is the distance, measured in units of 0.1 inch (2.54 mm), from the left edge of the document to the left boundary of the first field. The limiting range of this position is 4 to 85.

IMAGE= {YES|NO}: This operand specifies whether the data record should be in standard mode (IMAGE=NO), or image mode (IMAGE=YES). If this operand is not specified, IMAGE=NO is assumed.

NOSCAN=(field-end,...): Specifies an area on the document line that is to be ignored by the 3886. The parameter **field-end** is a decimal number indicating the distance, measured in units of 0.1 inch (2.54 mm), from the left edge of the document to the right end of the NOSCAN field. The field immediately to the left of the NOSCAN field must end with an address delimiter rather than a character delimiter.

Field Information Entries

**FLDn=({address-delimiter|character-delimiter}
,[field-length][, {NCRIT|font-code|
NCRIT,font-code}]):**

Describes each of the fields in a line. The n suffix is a number from 1 through 14 and the parameters are the same for keywords FLD1 through FLD14. The following rules apply when specifying these keywords:

- Fields may be described in any order in the macro.
- Each EDITn parameter must follow its associated FLDn parameter.
- The n suffix need not be 1 for the first field in the line; however, the n suffix must increase for each field from left to right on the document line.

address-delimiter is a decimal number that specifies the distance, measured in units of 0.1 inch (2.54 mm), from the left edge of the document to the right end of the field being defined. The last field in a line must end with an address delimiter.

character delimiter specifies the character that indicates the end of a field. The character delimiter is not considered part of the data; it is not included in the data record nor used in determining the length of the field.

Apostrophes enclosing the characters are optional for all characters except 0 through 9, and the special characters used in macro operands. For these characters, the apostrophes are required. For a description of these characters, see *OS/VS-DOS/VSE-VM/370 Assembler Language*, as listed in the Preface.

If a field ends with a character delimiter, the next field must be read using a font from the same font group. The font groups are:

- NPH1, NPH2, GOTH
- ANA1, ANA2, NUMA, MRKA
- NUMB, MRKB
- ANB1

field-length is a decimal number specifying the length of the field in the edited record. The length specified cannot be less than 1 or more than 127. If IMAGE=NO is specified, this parameter is required; if IMAGE=YES is specified, this parameter is invalid. The length specified in this parameter refers to the length of the field after any EDITn options have been performed. The sum of the field lengths for a line cannot be greater than 130.

NCRIT indicates that this is not a critical field. If this parameter is omitted, the field is assumed to be critical.

font-code specifies a font for this field, different from the font specified in the DFR macro. If this parameter is not specified, the font specified in the DFR macro is used for the field. For information about the valid codes, see the DFR macro description.

EDITn=({code|EDCHAR|code,EDCHAR}):
Describes the editing functions to be performed on the data by the 3886.

The parameters are the same for keywords EDIT1 through EDIT14. There must be a FLDn keyword corresponding with each EDITn keyword you specify. If an EDITn keyword is specified, a code, EDCHAR, or both must be specified. When image mode is used, the EDITn keywords are invalid.

When the editing functions are completed and the field is greater than the specified length, the field is truncated from the right and the wrong length field indicator is set on in the header record. If only blanks are truncated, the wrong length field indicator is not set.

code specifies the blanks to be removed and the fill characters to be added to the field, if any. The valid codes and their meanings are:

Code	Meaning
HLBLOF	All high- and low-order blanks are removed, the data is left justified, and the field is padded with blanks on the right (see Note).
ALBLOF	All blanks are removed from the data, the data is left-justified, and the field is padded with blanks on the right.
NOBLOF	No blanks are removed, the data is left-justified, and the field is padded on the right with blanks.
HLBHIF	All high- and low-order blanks are removed, the data is right-justified, and the field is padded to the left with EBCDIC zeros (X'F0') (see Note).
ALBHIF	All blanks are removed, the data is right-justified, and the field is padded with EBCDIC zeros (X'F0') on the left.
ALBNOF	All blanks are removed; the data must be equal in length to the field length specified. No padding is done.

Note: Two consecutive embedded blanks is the maximum number sent.

If the EDITn keyword is omitted or if EDITn is specified and the code is omitted, ALBLOF is assumed.

EDCHAR indicates that the characters specified in the EDCHAR keyword of the DFR macro are to be deleted from the field. If this parameter is omitted, the characters are not deleted.

FREND={YES|NO}: Indicates whether this is the last DLINT macro for the format record. NO indicates that more DLINT macros follow; YES indicates that this is the last one. If this operand is omitted, NO is assumed.

DTFDU Macro

The DTFDU macro defines sequential (consecutive) processing for a file contained on a diskette.

Applies to				
Input	Output			
x		M	EOFADDR=xxxxxxxx	Name of your end-of-file routine. (Required for input only).
x	x	M	IOAREA1=xxxxxxxx	Name of first I/O area.
x	x	M	RECSIZE=nnn	Length of one record in bytes.
x	x	O	CMDCHN=nn	Number of read/write CCWs (records) to be command-chained.
x	x	O	DEVADDR=SYSxxx	Symbolic unit, required only when not provided on an EXTENT statement.
x	x	O	DEVICE=3540	Must be 3540. If omitted, 3540 is assumed.
x	x	O	ERREXT=YES	Indicates additional errors and ERET desired. Specify ERROPT.
x	x	O	ERROPT=xxxxxxxx	IGNORE, or SKIP, or name of error routine.
x	x	O	FEED=xxx	YES means feed at end-of-file. NO means no feed, YES assumed if omitted.
	x	O	FILESEC=YES	YES means create file secure.
x	x	O	IOAREA2=xxxxxxxx	Name of second I/O area, if two areas are used.
x	x	O	IOREG=(nn)	Register number. General register 2 to 12 in parentheses. Omit WORKA.
x	x	O	MODNAME=xxxxxxxx	Name of DUMODFx logic module for this DTF. If omitted, IOCS generates standard name.
x	x	O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module.
x	x	O	SEPASMB=YES	DTFDU is to be assembled separately.
x	x	O	TYPEFLE=xxxxxx	INPUT or OUTPUT. If omitted, INPUT is assumed.
x		O	VERIFY=YES	3741/3742 input is verified.
x		O	VOLSEQ=YES	YES means OPEN is to check sequencing of multi-volume files.
x	x	O	WORKA=YES	GET or PUT specifies work area. Omit IOREG.
	x	O	WRTPROT=YES	File will be created with Write-Protect on (cannot be overwritten).

M=Mandatory
O=Optional

Figure 3-17. DTFDU macro operands.

CMDCHN=nn: This operand is specified to indicate the number of Read/Write CCWs to be command chained. Valid entries are 1, 2, 13, or 26; 1 is assumed if this operand is omitted. For each CCW specified by this operand, one record is processed (for example, if you code CMDCHN=13, 13 records are command chained and are processed – read or written – as a group). For entries of 2, 13, or 26,

either the IOREG operand or the WORKA operand must be specified.

DEVADDR=SYSxxx: This operand specifies the symbolic unit (SYSxxx) associated with the file if an EXTENT job control statement is not provided. An EXTENT statement is not required for single-volume input files. If an EXTENT statement is provided, its specification overrides any DEVADDR specification.

SYSxxx represents an actual I/O device address, and is used in the ASSGN job control statement to assign the actual I/O device address to this file.

DEVICE=3540: This operand specifies that the file to be processed is on the 3540. This operand may be omitted.

EOFADDR=name: This operand specifies the symbolic name of your end-of-file routine. IOCS automatically branches to this routine on an end-of-file condition. You can perform any operations required for the end-of-file in this routine (you will generally issue the CLOSE or CLOSER macro).

ERREXT=YES: This operand enables your ERROPT routine to return to DUMODFx with the ERET macro. It also enables permanent errors to be indicated to your program. For ERREXT facilities, the EROPT operand must be specified. However, to take full advantage of this option, use the ERROPT=name operand.

ERROPT= {IGNORE|SKIP|name}: Specify this operand if you do not want a job to be terminated when a permanent error cannot be corrected in the diskette error routine. If attempts to reread a chain of records are unsuccessful, the job is terminated unless the ERROPT entry is included. Either IGNORE, SKIP, or the name of an error routine can be specified. The functions of these parameters are described below.

IGNORE

> The error condition is ignored. The records are made available for processing. On output, the error condition is ignored and the records are considered written correctly.

SKIP

> No records in the error chain are made available for processing. The next chain of records is read from the diskette, and processing continues with the first record of that chain. On output, the SKIP option is the same as the IGNORE option.

name

> IOCS branches to your error routine named by this parameter regardless of whether or not ERREXT=YES is specified. In this routine you can process or make note of the error condition as desired.

If ERREXT is not specified, register 1 contains the address of the first record in the error chain. When processing in the ERROPT routine, reference records in the error chain by referring to the address supplied in register 1. The contents of the IOREG register or work area are variable and should not be used

to process error records. Also, GET macros must not be issued for records in the error chain. If any other IOCS macros (excluding ERET if ERREXT=YES) are used in this routine, the contents of register 13 (with RDONLY) and 14 must be saved and restored after their use. At the end of the routine, return control to IOCS by branching to the address in register 14. For a read error, IOCS skips that error chain of records, and makes the first record of the next chain available for processing in the main program.

If ERREXT is specified, register 1 contains the address of a two part parameter list containing the 4-byte DTFDU address and the 4-byte address of the first record in the error chain. Register 14 contains the return address. Processing is similar to that described above except for addressing the records in error.

At the end of its processing, the routine returns to LIOCS by issuing the ERET macro.

For an input file, the program:
- skips the error chain and reads the next chain with an ERET SKIP,
- ignores the error with an ERET IGNORE,
- it makes another attempt to read the error chain with an ERET RETRY.

For an output file the only acceptable parameters are IGNORE or name, and the program
- ignores the error condition with ERET IGNORE or ERET SKIP,
- attempts to write the error chain with an ERET RETRY. Bad spot control records (1, 2, 13, or 26 records depending on the CMDCHN specification) are written at the current diskette address, and the write chain is retried in the next 1, 2, 13, or 26 (depending on the CMDCHN specification) sectors on the disk.

The DTFDU error options are shown in Figure 3-18.

To Terminate the job,	specify nothing;
Skip the error record,	specify ERROPT=SKIP;
Ignore the error record,	specify ERROPT=IGNORE;
Process the error record,	specify ERROPT=name;
after processing the record, to leave the error-processing routine and	
Skip the (input) record,	execute ERET SKIP;
Ignore the record,	execute ERET IGNORE;
Retry reading or writing the record,	execute ERET RETRY.

Figure 3-18. DTFDU error options.

FEED= {YES|NO}: If YES is specified and IOCS detects an end-of-file condition, the diskette being processed is fed to the stacker and a new diskette is fed to the disk drive (providing another diskette is

still in the hopper). If NO is specified, the diskette is left mounted for the next job. If the operand is omitted, YES is assumed.

FILESEC=YES: This operand applies to output only. On output it causes OPEN or OPENR to set the security flag in the file label. For subsequent input, the security flag causes an operator message to be written. The operator must then reply in order to make the file available to be read.

Note: When this operand is used with WRTPROT=YES, the reuse of the diskette is prevented.

IOAREA1=name: This operand specifies the symbolic name of the I/O area used by the file. IOCS either reads or writes records using this area. Note that you should provide an I/O area equal in size to the result obtained from multiplying the RECSIZE entry by the CMDCHN entry.

IOAREA2=name: If two I/O areas are used by GET or PUT, this operand is specified. You should provide an I/O area equal in size to the result obtained from multiplying the RECSIZE entry by the CMDCHN entry.

IOREG=(r): This operand specifies the general purpose register (one of 2 to 12) in which IOCS puts the address of the logical record that is available for processing. At OPEN or OPENR time, for output files, IOCS puts the address of the area where the user can build a record in this register. The same register can be used for two or more files in the same program, if desired. If this is done, the problem program must store the address supplied by IOCS for each record. If this operand is specified, omit the WORKA operand.

This operand must be specified if the CMDCHN factor is 2 or higher and records are processed in one I/O area, or if two I/O areas are used and records are processed in both I/O areas.

MODNAME=name: This operand specifies the name of the logic module which is to process the file. If the logic module is assembled with the program, MODNAME must specify the same name as the DUMODFx macro. If this operand is omitted, standard names are generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte double-word aligned save area. Each task should have its own uniquely defined save area. When an imperative macro (except OPEN, OPENR) is issued, register 13 must contain the address of the save area associated with the task. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If an ERROPT routine issues I/O macros using the same read-only module that caused control to pass to the error routine, your problem program must provide another save area. One save area is used for the normal I/O operations, and the second for input/output operations in the ERROPT routine. Before returning to the module that entered the ERROPT routine, register 13 must be set to the save area address originally specified for that DTF.

If this operand is omitted, the generated module is not reentrant and no save area need be established.

RECSIZE=nnn: This operand specifies (in bytes) the length of each record in the input/output area (1 to 128 bytes).

SEPASMB=YES: Include this operand only if the DTFDU is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an entry point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TYPEFLE= {INPUT|OUTPUT}: This operand indicates whether the file is an input or output file.

VERIFY=YES: This operand specifies that the input on a 3741/3742 must be verified before processing may continue. If VERIFY=YES is not specified, it is assumed that the input need not be verified. If VERIFY=YES is specified and the input is not verified, the job is canceled and message 4n57I is issued. If the operand is specified for an output file, it will be ignored.

VOLSEQ=YES: This operand is only valid on input. If specified, it causes OPEN or OPENR to ensure that the volume sequence numbers of a multi-volume file are in ascending and sequential order. However, if the volume sequence number of the first volume processed is blank, no volume sequence checking is done.

WORKA=YES: If I/O records are processed or built in work areas instead of in the I/O areas, specify this operand. You must set up the work area in

storage. The address of the work area, or a general register containing the address, must be specified in each GET or PUT macro. For a GET or PUT macro, IOCS moves the record to or from the specified work area.

When this operand is specified, the IOREG operand must be omitted.

WRTPROT=YES: This operand indicates that an output file will be created with Write-Protect (meaning that the file cannot be overwritten). For 3540 support, this has no effect on subsequent input processing of the file.

Note: When this operand is used with FILESEC=YES, the reuse of the diskette is prevented.

DUMODFx Macro

Two categories of file characteristics are defined for diskette unit module generation macros:

- DUMODFI - Diskette Unit MODule, Fixed length records, Input file.
- DUMODFO - Diskette Unit MODule, Fixed length records, Output file.

The macro operation and the keyword operands define the characteristics of the module.

A module name can be contained in the name field of this macro. The macro operation is contained in the operation field, either DUMODFI (for input) or DUMODFO (for output).

ERREXT=YES: Include this operand if permanent errors are returned to a problem program ERROPT routine or if the ERET macro is used with the DTF and module. The ERROPT operand must be specified for this module.

ERROPT=YES: This operand applies to both DUMODFX macros. This operand is included if the module handles any of the error options for an error chain. Logic is generated to handle any of the three options (IGNORE, SKIP, or name) regardless of which option is specified in the DTF. This module also processes any DTF in which the ERROPT operand is not specified.

If this operand is not included, your program is canceled whenever a permanent error is encountered.

RDONLY=YES: This operand causes a read-only module to be generated. If this operand is specified, any DTF used with this module must have the same operand.

SEPASMB=YES: Include this operand only if the logic module will be assembled separately. This causes a CATALR card with the module name

(standard or user-specified) to be punched ahead of the object deck, and defines the module name to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the logic module is being assembled with the problem program and no CATALR is punched.

Standard DUMOD Names

Each name begins with a 3-character prefix (IJN) and continues with a 5-character field corresponding to the options permitted in the generation of the module, as shown below.

DUMODFx name = IJNabcde

a = D

b = I DUMODFI
   = O DUMODFO

c = C ERROPT=YES and ERREXT=YES
   = E ERROPT=YES
   = Z neither is specified

d = Z

e = Y RDONLY=YES
   = Z RDONLY not specified

Subset/Superset DUMOD Names

Figure 3-19 illustrates the subsetting and supersetting allowed for DUMOD names.

```
                * + * *
    I J N D  I  C Z Y
              O E   Z
              Z
```

+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.
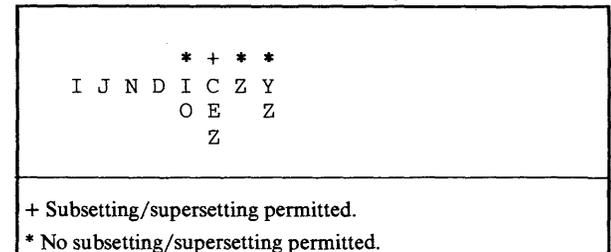
Figure 3-19. Subsetting and supersetting of DUMOD names.

DTFIS Macro

The DTFIS macro defines a DASD file for the Indexed Sequential Access Method.

Applies to						
Ran. Rtvl.	Seq. Rtvl.	Load	Add			
x	x	x	x	M	DSKXTNT=n	Maximum number of extents specified for this file
x	x	x	x	M	IOROUT=xxxxxx	(LOAD, ADD, RETRVE, or ADDRTR)
x	x	x	x	M	KEYLEN=nnn	Number of bytes in record key (maximum is 255)
x	x	x	x	M	NRECDS=nnn	Number of records in a block. Specify for blocked records only; if unblocked, 1 is assumed.
x	x	x	x	M	RECFORM=xxxxxx	(FIXUNB or FIXBLK)
x	x	x	x	M	RECSIZE=nnnn	Number of characters in logical record.
x	x	x	x	O	CYLOFL=nn	Number of tracks for each cylinder overflow area. Maximum = 8 for 2311, 18 for 2314 , 17 for 3330 and 3333, 10 for 3340
x	x	x	x	O	DEVICE=nnnn	(2311, 2314, 3330, 3340). If omitted, 2311 is assumed.
x	x	x	x	O	ERREXT=YES	Non data-transfer error returns and ERET desired.
x	x	x	x	O	HINDEX=nnnn	(2311, 2314, 3330, 3340). Unit containing highest level index. If omitted, 2311 is assumed.
x	x		x	O	HOLD=YES	Track hold function is desired
x			x	O	INDAREA=xxxxxxxx	Symbolic name of cylinder index area
x			x	O	INDSKIP=YES	Index skip feature is to be used
x			x	O	INDSIZE=nnnnn	Number of bytes required for the cylinder index area
		x	x	O	IOAREAL=xxxxxxx	Name of I/O area
x				O	IOAREAR=xxxxxxx	
	x			O	IOAREAS=xxxxxxx	
	x	x		O	IOAREA2=xxxxxxx	Name of second I/O area
x	x			O	IOREG =(nn)	Register number. Omit if WORKA or WORKS is specified
			x	O	IOSIZE=nnnn	Bytes alloted to IOAREAL
x	x			O	KEYARG=xxxxxxx	Name of key field in storage, for random retrieval or sequential retrieval starting by key
x	x	x	x	O	KEYLOC=nnnn	Number of high-order position of key field within record, if RECFORM=FIXBLK
x	x	x	x	O	MODNAME=xxxxxxxx	Name of ISMOD logic module for this DTF. If omitted, IOCS generates standard name
x	x	x	x	O	MSTIND=YES	Master index used.
x	x	x	x	O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module.
x	x	x	x	O	SEPASMB=YES	DTFIS is to be assembled separately.
x	x			O	TYPEFLE=xxxxxx	(RANDOM, SEQNTL, or RANSEQ)
x	x	x	x	O	VERIFY=YES	Check disk records after they are written.
		x	x	O	WORKL=xxxxxxxx	Name of work area for loading or adding to the file.
x				O	WORKR=xxxxxxxx	Name of work area for random retrieval. Omit IOREG
	x			O	WORKS=YES	GET or PUT specifies work area

M=Mandatory
O=Optional

Figure 3-20. DTFIS macro.

CYLOFL=n: This operand must be included if cylinder overflow areas are reserved for a file. Do not include this entry if no overflow areas are reserved.

When a file is loaded or when records are added, this operand is required to reserve the areas for cylinder overflow. It specifies the number of tracks to be reserved on each cylinder. The maximum num-

ber of tracks that can be reserved on each cylinder is:

for 2311	8
for 2314, or 2319	18
for 3330 or 3333	17
for 3340	10

DEVICE= {2311|2314|3330|3340}: This operand specifies the unit that contains the prime data area or overflow areas for the logical file. For ISAM the prime data area must be on the same device type, and for a 3340 on the same model of data module.

For devices supported by DOS/VSE and not included in the above operand specification, specify device codes as listed in Figure 3-21.

DSKXTNT=n: This operand must be included to specify the maximum number of extents for this file. The number must include all the data area extents if more than one DASD area is used for the data records, and all the index area and independent overflow area extents that are specified by EXTENT job control statements. Thus the minimum number specified by this entry is 2: one extent for one prime data area, and one for a cylinder index. Each area assigned to an ISAM file is considered an extent.

Note: Master and cylinder indexes are treated as one area. When there is one master index extent, one cylinder index extent, and one prime data area extent, DSKXTNT=2 could be specified.

ERREXT=YES: This operand is required for IOCS to supply your program with detailed information about unrecoverable I/O errors occurring before a data transfer takes place, and for your program to be able to use the ERET imperative macro to return to IOCS specifying an action to be taken for an error condition.

Some error information is available for testing by your program after each imperative macro is executed, even if ERREXT=YES is not specified, by referencing field *filenameC*. Filename is the same name as that specified in the DTF header entry for the file. One or more of the bits in the filenameC byte may

be set to 1 by IOCS. The meaning of the bits varies depending on which parameter was specified in the IOROUT operand; Figure 3-22 shows the meaning if IOROUT=ADD, RETRVE, or ADDRTR was specified; Figure 3-23 shows the meaning if IOROUT=LOAD was specified.

If ERREXT=YES is not specified, IOCS returns the address of the DTF table in register 1, as well as any data-transfer error information in filenameC, after each imperative macro is executed; non-data-transfer error information is not given. After testing filenameC, return to IOCS by issuing any imperative macro except ERET; no special action is taken by IOCS to correct or check an error.

If ERREXT=YES is specified, IOCS returns the address of an ERREXT parameter list in register 1 after each imperative macro is executed, and information about both data-transfer and non-data-transfer errors in filenameC. The format of the ERREXT parameter list is shown in Figure 3-25. After testing filenameC and finding an error, return to IOCS by using the ERET imperative macro; IOCS takes the action indicated by the ERET operand. If HOLD=YES (and ERREXT=YES), ERET must be used to return to IOCS to free any held track.

In your program, you should check byte 16, bit 7 of the DTF for a blocksize compatibility error when adding to, or extending a file. If the blocksize of your program is not equal to the blocksize of the previously built file, this bit will be set to 1.

HINDEX= {2311|2314|3330|3340}: This entry specifies the unit containing the highest index.

For devices supported by DOS/VSE and not included in the above operand specification, specify device codes as listed in Figure 3-24.

Placing the highest index on a separate unit is recommended only if that unit is physically separate from the unit(s) holding the track indexes and the data of the file, and if it has its own access mechanism.

DEVICE = specification	Device in use					
	2311	2314	2319	3330-1,2*	3340, 35MB	3340, 70MB
Default	x					
2311	x					
2314		x	x			
3330				x		
3340					x	x

* Also 3350 in 3330-1 compatibility mode.

Figure 3-21. DEVICE= specifications for DTFIS.

Bit	Cause	Explanation
0	DASD error	An uncorrectable DASD error has occurred (except wrong length record).
1	Wrong length record	A wrong length record has been detected during an I/O operation.
2	End of file	The EOF condition has been encountered during execution of the sequential retrieval function.
3	No record found	The record to be retrieved has not been found in the file. This applies to Random (RANSEQ) and to SETL in SEQNTL (RANSEQ) when KEY is specified, or after GKEY.
4	Illegal ID specified	The ID specified to the SETL in SEQNTL (RANSEQ) is outside the prime file limits.
5	Duplicate record	The record to be added to the file has a duplicate record key of another record in the file.
6	Overflow area full	An overflow area in a cylinder is full, and no independent overflow area has been specified; or an independent overflow area is full, and the addition cannot be made. You should assign an independent overflow area or extend the limit.
7	Overflow	The record being processed in one of the retrieval functions (RANDOM/SEQNTL) is an overflow record.

Figure 3-22. FilenameC - status or condition code byte if IOROUT=ADD, RETRVE, or ADDRTR

Bit	Cause	Explanation
0	DASD error	An uncorrectable DASD error has occurred (except wrong length record).
1	Wrong length record	A wrong length record has been detected during an I/O operation.
2	Prime area full	The next to the last track of the prime data area has been filled during the load or extension of the file. You should issue the ENDFL macro, then do a load extend on the file with new extents given.
3	Cylinder index area full	The cylinder index area is not large enough to contain all entries needed to index each cylinder specified for the prime data area. This condition can occur during the execution of the SETFL. You must extend the upper limit of the cylinder index by using a new extent card.
4	Master index full	The master index area is not large enough to contain all the entries needed to index each track of the cylinder index. This condition can occur during SETFL. You must extend the upper limit, if you are creating the file, by using an extent card. Or, you must reorganize the file and assign a larger area.
5	Duplicate record	The record being loaded is a duplicate of the previous record.
6	Sequence check	The record being loaded is not in the sequential order required for loading.
7	Prime data area overflow	There is not enough space in the prime data area to write an EOF record. This condition can occur during the execution of the ENDFL macro.

Figure 3-23. FilenameC - status or condition code byte if IOROUT=LOAD.

HOLD=YES: This operand provides for the track hold option for both data and index records. If the HOLD operand is omitted, the track hold function is not performed. Because track hold cannot be performed on a LOAD file, HOLD=YES cannot be specified when IOROUT=LOAD.

If HOLD=YES and ERREXT=YES, your program must issue the ERET macro to return to the ISAM module to free any held tracks.

INDAREA=name: This operand specifies the name of the area assigned to the cylinder index. If specified, all or part of the cylinder index resides in virtual storage thereby increasing throughput. If this operand is included, INDSIZE must be included.

If the area assigned to INDAREA is large enough for all the index entries to be read into virtual storage at one time and the index skip feature (INDSKIP) is not specified, no presorting of records need be done. If the area assigned to INDAREA is not large enough, the records processed should be presorted to fully utilize the resident cylinder index.

INDSKIP=YES: When cylinder index entries reside in virtual storage, this operand specifies the index skip feature. This feature allows ISAM to skip any index entries preceding those needed to process

HINDEX = specification	Device in use					
	2311	2314	2319	3330-1,2*	3340,35MB	3340,70MB
Default	x					
2311	x					
2314		x	x			
3330				x		
3340					x	x

* Also 3350 in 3330-1 compatibility mode.

Figure 3-24. HINDEX= specifications for DTFIS.

Bytes	Bits	Contents
0-3	--	DTF address
4-7	--	Virtual storage address of the record in error
8-15	--	DASD address (mbbcchhr) of the error where m is the extent sequence number and r is a record number which can be inaccurate if a read error occurred during a read of the highest level index.
16	.	Record identification:
	1	Data record
	2	Track index record
	3	Cylinder index record
		Master index record
	.	Type of operation:
	4	Not used
	5	Not used
	6	Read
	7	Write
17	--	Command code of failing CCW

Figure 3-25. ERREXT parameter list.

a given key. If the index skip operand is omitted, the cylinder indexes are processed sequentially.

This operand may be specified only with the INDAREA and INDSIZE operands and increases throughput only when:

- The records are presorted.
- The allocated virtual storage is insufficient for storing all of the cylinder index.
- One or more large segments of the file are not referenced.

INDSIZE=n: This operand specifies the length (in bytes) of the index area assigned in virtual storage to the cylinder index by INDAREA. The minimum you can specify is:

$$n=(m+3)(keylength+6)$$

where

m = the number of entries to be read into virtual storage at a time.

3 = the number of dummy entries

6 = a pointer to the cylinder

If m is set equal to the number of prime data cylinders+1, the entire cylinder index is read into virtual storage at one time. The maximum value for n = 32767.

The resident index facility is suppressed if this operand is omitted, the minimum requirement is not met at assembly time, or an unrecoverable read error is encountered while reading the index.

IOAREAL=name: This operand must be included when a file is created (loaded) or when records are added to a file. It specifies the name of the output area used for loading or adding records to the file. The specified name must be the same as the name used in the DS instruction that reserves the area of storage. The ISAM routines construct the contents of this area and transfer records to DASD.

This output area must be large enough to contain the count, key, and data areas of records. Furthermore, the data-area portion must provide enough space for the sequence-link field of overflow records whenever records are added to a file (see Figure 3-26).

If IOAREAL is increased to permit the reading and writing of more than one physical record on DASD at a time, the IOSIZE operand must be included when records are added to the file. In this case, the IOREAL area must be at least as large as the number of bytes specified in the IOSIZE operand.

When simultaneously building two ISAM files using two DTFs, do not use a common IOAREAL. Also, do not use a common area for IOAREAL, IOAREAR, and IOAREAS in multiple DTFs.

IOAREAR=name: This operand must be included whenever records are processed in random order. It specifies the name of the input/output area for random retrieval (and updating). The specified name must be the same as that used in the DS instruction that reserves this area of storage.

The I/O area must be large enough to contain the data area for records. Furthermore, the data-area

FUNCTION	OUTPUT AREA REQUIREMENTS (IN BYTES)			
	Count	Key	Sequence Link	Data
Load Unblocked Records	8	Key Length	—	Record Length
Load Blocked Records	8	Key Length	—	Record Length x Blocking Factor
Add Unblocked Records	8	Key Length	10	Record Length
Add Blocked Records	8	Key Length	—	Record Length x Blocking Factor
	8	Key Length	10	Record Length
* Whichever Is Larger				

OR* (between the two Add Blocked Records rows)

Figure 3-26. Output area requirements for loading or adding records to a file by ISAM.

FUNCTION	I/O AREA REQUIREMENTS (IN BYTES)			
	Count	Key	Sequence Link	Data
Retrieve Unblocked Records	—	Key Length for sequential unblocked records	10	Record Length
Retrieve Blocked Records	—	——	—	Record Length (including keys) x Blocking Factor
	—	——	10	Record Length
* Whichever Is Larger				

OR* (between the two Retrieve Blocked Records rows)

Figure 3-27. I/O area requirements for random or sequential retrieval by ISAM.

portion must provide enough space for the sequence-link field of overflow records (see Figure 3-27).

IOAREAS=name: This operand must be included whenever records are processed in sequential order by key. It specifies the name of the input/output area used for sequential retrieval (and updating). The specified name must be the same as that used in the DS instruction that reserves this area of storage.

This I/O area must be large enough to contain the key and data areas of unblocked records and the data area for blocked records. Furthermore, the data-area portion must provide enough space for the sequence-link field of overflow records (see Figure 3-27).

IOAREA2=name: This operand permits overlapping of I/O with indexed sequential processing for either the load (creation) or sequential retrieval functions. Specify the name of an I/O area to be used when loading or sequentially retrieving records. The I/O area must be at least the length of the area specified by either the IOAREAL operand for the load function or the IOAREAS operand for the sequential retrieval function. If the operand is omitted, one I/O area is assumed. If TYPEFLE=RANSEQ. this operand must not be specified.

IOREG=(r): This operand must be included whenever records are retrieved and processed directly in the I/O area. It specifies the register that ISAM uses to indicate which individual record is available for processing. ISAM puts the address of the current record in the designated register (any of 2 through 12) each time a READ, WRITE, GET, or PUT is executed.

IOROUT= {LOAD|ADD|RETRVE|ADDRTR}: This entry must be included to specify the type of function to be performed. The parameters have the following meanings:

LOAD
To build a logical file on a DASD or to extent a file beyond the highest record presently in a file.

ADD
To insert new records into a file.

RETRVE
To retrieve records from a file for either random or sequential processing and/or updating

ADDRTR
To both insert new records into a file (ADD) and retrieve records for processing and/or updating (RTR).

IOSIZE=n: This operand specifies the (decimal) number of bytes in the virtual-storage area assigned for the add function using IOAREAL. The number n can be computed using the following formula:

$$n = m(keylength+blocksize+40)+24$$

where m is the maximum number of physical records that can be read into virtual storage at one time; 40 is the sum of 8 for the count field and 32 for an ISAM CCW; 24 is another ISAM CCW. The number n must be at least equal to

$$(keylength+blocksize+74)$$

This formula accounts for a needed sequence link field for unblocked records or short blocks (see Figure 3-26 and Figure 3-27).

If the operand is omitted, or if the minimum requirement is not met, no increase in throughput is realized.

The number n should not exceed the track capacity because throughput cannot be increased by specifying a number larger than the capacity of a track.

KEYARG=name: This operand must be included for random READ/WRITE operations and sequential retrieval initiated by key. It specifies the symbolic name of the key field in which you must supply the record key to ISAM.

KEYLEN=n: This operand must be included to specify the number of bytes in the record key.

KEYLOC=n: This operand must always be specified if RECFORM=FIXBLK. It supplies ISAM with the high-order position of the key field within the data record. That is, if the key is recorded in positions 21-25 of each record in the file, this operand should specify 21.

ISAM uses this specification to locate (by key) a specified record within a block. The key area of a block of records contains the key of the highest record in the block. To search for any other records, ISAM locates the proper block and then examines the key field within each record in the block.

MODNAME=name: This operand may be used to specify the name of the logic module used with the DTF table to process the file. If the logic module is assembled with the program, the MODNAME in the DTF must specify the same name as the ISMOD macro. If this entry is omitted, standard names are

generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called.

MSTIND=YES: This operand is included whenever a master index is used or is to be built for a file. The location of the master index is specified by an EXTENT job control statement.

NRECDS=n: This operand specifies the number of logical records in a block (called the blocking factor). It is required only if RECFORM=FIXBLK. For FIXBLK, n must be >1; for FIXUNB, n must be $=1$.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each task should have its own uniquely defined save area. Register 13 must contain the address of the save area associated with the task each time an imperative macro (except OPEN, OPENR, LBRET, SETL, or SETFL) is issued. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

RECFORM= {FIXUNB|FIXBLK}: This operand specifies whether records are blocked or unblocked. FIXUNB is used for unblocked records, and FIXBLK for blocked records. If FIXBLK is specified, the key of the highest record in the block becomes the key for the block and must be recorded in the key area.

The specification that is included when the logical file is loaded onto a DASD must also be included whenever the file is processed.

Records in the overflow area(s) are always unblocked, but this has no effect on this operand. RECFORM refers to records in the prime data area only.

RECSIZE=n: This operand must be included to specify the number of characters in the data area of each individual record. This operand should specify the same number for additions and retrieval as indicated when the file was created.

SEPASMB=YES: Include this operand only if the DTFIS is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and defines the filename as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TYPEFLE= {RANDOM|SEQNTL|RANSEQ}:
This operand must be included when
IOROUT=RETRVE or IOROUT=ADDRTR. The ope-
rand specifies the type(s) of processing performed by
your program for the file.

RANDOM
> is used for random processing. Records are
> retrieved in random order specified by key.

SEQNTL
> is used for sequential processing. Your pro-
> gram specifies the first record retrieved, and
> thereafter ISAM retrieves records in sequential
> order by key. The first record is specified by
> key, ID, or the beginning of the logical file see
> "SETL Macro".

RANSEQ
> is used if both random and sequential process-
> ing are to be performed for the same file. If
> RANSEQ is specified, the IOAREA2 operand
> must not be specified.

TYPEFLE is not required for loading or adding
functions.

VERIFY=YES: Use this operand if you want to
check the parity of disk records after they are writ-
ten. If this operand is omitted, any records written
on a disk are not verified.

WORKL=name: This operand must be included
whenever a file is created (loaded) or records are
added to a file. It specifies the name of the work
area in which you must supply the data records to
ISAM for loading or adding to the file. The specified
name must be the same as the name used in the DS
instruction that reserves this area of storage.

This work area must provide space for one logical
record when a file is created (for blocked records:
data; for unblocked records: key and data).

The original contents of WORKL are changed due
to record shifting in the ADD function.

WORKR=name: When records are processed in
random order, this operand must be included if the
individual records are to be processed in a work area
rather than in the I/O area. It specifies the name of
the work area. This name must be the same as the
name used in the DS instruction that reserves this
area of storage. This area must provide space for
one logical record (data area). When this entry is
included and a READ (or WRITE) macro is executed,
ISAM moves the individual record to (or from) this
area.

WORKS=YES: When records are processed in
sequential order, this operand must be included if
the individual records are processed in work areas
rather than in the I/O area. Each GET and PUT ma-
cro must specify the name of the work area to or
from which ISAM is to move the record. When proc-
essing unblocked records, the area must be large
enough for one record (data area) and the record
key (key area). For blocked records, the area must
be large enough for one logical record (data area)
only. The work area requirements are as shown in
Figure 3-28.

	Unblocked Records	Blocked Records
Load	(KL + DL) or 10*	DL or 10*
ADD	(KL + DL) or 10*	DL or (KL + 10)*
Random Retrieve	DL	DL
Sequential Re-trieve	KL + DL	DL
K=KEY, D=Data, L=Length * Whichever is greater		

Figure 3-28. Work area requirements.

ISMOD Macro

Operand	Remarks
ERREXT=YES[1]	Required if non-data-transfer error conditions or ERET are desired.
CORDATA=YES	Required to add records using the DTF IOSIZE operand.
CORINDX=YES[2]	Required to add or retrieve records with the cylinder index entries in virtual storage.
HOLD=YES	Specifies the track hold option.
IOAREA2=YES	Required if two I/O areas are to be used.
IOROUT= {LOAD| ADD| RETRVE[2]| ADDRTR}	Specifies function to be performed.
RDONLY=YES[1]	Required if a read-only module is to be generated.
RECFORM= {FIXUNB| FIXBLK| BOTH[1]}	Describes file. Required if IOROUT specifies ADD or ADDRTR. If IOROUT specifies LOAD or RETRVE, BOTH is assumed.
RPS=SVA	To assemble RPS logic modules.
SEPASMB=YES	If the module is assembled separately.
TYPEFLE= {RANDOM| SEQNTL| RANSEQ}	Required if IOROUT specifies RETRV or ADDRTR.

[1] Value assumed if RPS=SVA specified.
[2] Value assumed if RPS=SVA and IOROUTE=ADD or RETRVE specified.

Figure 3-29. Operands of the ISMOD macro.

Note: If an ISMOD module precedes an assembler language USING statement or follows your program, registers 2-12 remain unrestricted even at assembly time. However, if the ISMOD module lies within your program, you should issue the same USING statement (as that which was issued before the ISMOD module) directly following the module. This action is necessary because the ISMOD module uses registers 1, 2, and 3 as base registers, and the ISMOD CORDATA module uses registers 1, 2, 3, and 5 as base registers. Each time either module is assembled, these registers are dropped.

CORINDX=YES: Include this operand to generate a module that can process DTFIS files (add or random retrieve functions) with or without the cylinder index entries resident in virtual storage. If omitted, the module generated cannot process the resident cylinder index entries.

If an unrecoverable I/O error occurs while reading indexes into virtual storage, the program will not use the resident cylinder index entries.

CORDATA=YES: Include this operand if the mo-

dule is to add records to files with the IOSIZE DTFIS operand. If this operand is included, the IOSIZE operand is required in the DTF. If you omit the CORDATA=YES operand, you will not have an increase in throughput when adding records to a file.

ERREXT=YES: Include this operand if the ERET macro is to be used with this module or if non-data-transfer error conditions are returned in filenameC. If HOLD=YES and ERREXT=YES, your program must issue the ERET macro to return to the ISAM module to free any held tracks. See the DTF ERREXT and HOLD operands.

HOLD=YES: This operand provides for the track hold option for both data and index records. If the HOLD operand is omitted, the track hold function is not performed.

Because track hold cannot be performed on a LOAD file, HOLD=YES cannot be specified when IOROUT=LOAD.

If HOLD=YES and ERREXT=YES, your program must issue the ERET macro to return to the ISAM module to free any held tracks.

IOAREA2=YES: Include this operand if a second I/O area is to be used – that is, if IOAREA2 is specified in any of the DTFs linked to the logic module. The operand is only valid for load or sequential retrieval functions. The module can process DTFs with one or two I/O areas specified. This operand must not be specified if TYPEFLE=RANSEQ is specified.

IOROUT={LOAD\|ADD\|RETRVE\|ADDRTR}: This operand specifies the type of module required to perform a given function.

LOAD
    generates a module for creating or extending a file.

ADD
    generates a module for adding new records to an existing file.

RETRVE
    generates a module to retrieve, either randomly or sequentially, records from a file.

ADDRTR
    generates a module that combines the features of the ADD and RETRVE modules. This module also processes any file in which only ADD or RETRVE is specified in the IOROUT operand of the DTF, and in which the TYPEFLE operand contains the corresponding parameter (or a subset of it).

RDONLY=YES: This operand causes a read-only module to be generated. Whenever this operand is specified, any DTF used with the module must have the same operand.

RECFORM={FIXUNB|FIXBLK|BOTH}: This operand generates a module that creates, adds to, or processes an unblocked (FIXUNB) or blocked (FIXBLK) file. If BOTH is specified, a module is generated to process both unblocked and blocked files, and the DTF may specify either FIXUNB or FIXBLK in the RECFORM operand. The RECFORM operand is required only when IOROUT specifies ADD or ADDRTR. If IOROUT specifies LOAD or RETRVE, a module that handles fixed-length blocked and unblocked files is generated, and the operand is not required.

RPS=SVA This operand causes the RPS logic modules to be assembled.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and defines the module name as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

TYPEFLE={RANDOM|SEQNTL|RANSEQ}: This operand is required when IOROUT specifies RETRVE or ADDRTR. RANDOM generates a module that includes only random retrieval capabilities. SEQNTL generates a module that includes only sequential retrieval capabilities. RANSEQ generates a module that includes random and sequential capabilities. It also processes any file in which the TYPEFLE operand specifies either RANDOM or SEQNTL. If TYPEFLE=RANSEQ, IOAREA2=YES must not be specified.

When all operands are omitted, the ISMOD module can only process files where IOROUT=RETRVE, TYPEFLE=RANSEQ, CORINDX, CORDATA, HOLD, and RDONLY are not specified. The name of that module is IJHZRBZZ.

Standard ISMOD Names
Each name begins with a 3-character prefix (IJH) and continues with a 5-character field corresponding to the options permitted in the generation of the module.

ISMOD name = IJHabcde

a = A RECFORM=BOTH, IOROUT=ADD or ADDRTR
  = B RECFORM=FIXBLK, IOROUT=ADD or ADDRTR
  = U RECFORM=FIXUNB, IOROUT=ADD or ADDRTR
  = Z RECFORM is not specified. (IOROUT=LOAD or RETRVE)

b = A IOROUT=ADDRTR
  = I IOROUT=ADD
  = L IOROUT=LOAD
  = R IOROUT=RETRVE
  = V IOROUT=ADDRTR, RPS=SVA
  = X IOROUT=LOAD, RPS=SVA

c = B TYPEFLE=RANSEQ
  = G IOAREA2=YES, TYPEFLE=SEQNTL or IOROUT=LOAD
  = R TYPEFLE=RANDOM
  = S TYPEFLE=SEQNTL
  = Z neither is specified (IOROUT=LOAD or ADD)

d = B CORINDX=YES and HOLD=YES
  = C CORINDX=YES
  = O HOLD=YES
  = Z neither is specified

e = F CORDATA=YES, ERREXT=YES, RDONLY=YES
  = G CORDATA=YES and ERREXT=YES
  = O CORDATA=YES and RDONLY=YES
  = P CORDATA=YES
  = S ERREXT=YES and RDONLY=YES
  = T ERREXT=YES
  = Y RDONLY=YES
  = Z neither is specified

Subset/Superset ISMOD Names
Figure 3-30 shows the subsetting and supersetting allowed for ISMOD names. Five parameters allow supersetting. For example, the module IJHBABZZ is a superset of the module IJHBASZZ.

```
          +  +  +  +  +
   I J H  A  A  B  B  F
          B  I  R  O  O
          Z  +  +  +  +
          +  A  B  C  S
          A  R  S  Z  Y
          U  *  +     +
          Z  L  G     G
                S     P
                +     +
                G     T
                Z     Z
```

+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.

Figure 3-30. Subsetting and supersetting of ISMOD names.

If two or more modules with the same entry point are included, the linkage editor message 2143I,

(invalid duplication of entry point label) is generated. (Occasionally these entry points are not obvious when using the preceding chart, but the module can perform the indicated functions.) This message can usually be supressed by including a superset module. However, modules with and without prime data in main storage or modules with TYPEFLE=RANDOM and IOAREA2=YES cannot be combined. Therefore, you should take either of the following actions:

1. Specify prime data in core for each ADD type DTF in your program. In this case, superset modules are generated.
2. Specify the MODNAME operand in the DTF, and include an ISMOD of that name. The DTF then generates only the specified module.

DTFMR Macro

DTFMR defines an input file processed on a 1255, 1259, or 1419 magnetic character reader, or a 1270 or 1275 optical character reader/sorter.

M	DEVADDR=SYSnnn	Symbolic unit assigned to the magnetic character reader.
M	IOAREA1 = xxxxxxxx	Name of the document buffer area.
O	ADDAREA=nnn	Additional document buffer area (ADDAREA+RECSIZE=250). If omitted, no area is alloted.
O	ADDRESS=DUAL	Must be included only if the device is a 1419 or 1275 with a dual address adapter.
O	BUFFERS=nnn	Specifies the number of buffers needed. If omitted, 25 is assumed.
O	ERROPT=xxxxxxxx	Name of your error routine. Required only if the CHECK macro is used.
O	EXTADDR=xxxxxxxx	Name of your stacker selection routine. Required only if SORTMDE=ON.
O	IOREG=(nn)	Pointer register number. If omitted, register 2 is assumed. General registers 2-12, written in parentheses.
O	MODNAME=xxxxxxxx	Name of your I/O module. Required only if a nonstandard module is referenced.
O	RECSIZE=nnn	Specifies the maximum record length. If omitted, 80 is assumed.
O	SECADDR=SYSnnn	Specifies secondary symbolic unit assigned to (dual address) 1275 or 1419. Required only if LITE macro is used.
O	SEPASMB=YES	Required only if the DTF is assembled separately; otherwise it should be omitted.
O	SORTMDE=xxx	ON – 1255/1259/1270 or program sort mode used; OFF – 1419/1275 sort mode used. If omitted, ON is assumed.

M=Mandatory
O=Optional

Figure 3-31. DTFMR macro operands.

ADDAREA=n: This operand must be included only if an additional buffer work area is needed. The parameter n specifies the number of additional bytes you desire in each buffer. The sum of the ADDAREA and RECSIZE specifications must not exceed 250. This area can be used as a work area and/or output area and is reset to binary zeros when the next GET or READ for the file is executed.

ADDRESS=DUAL: This operand must be included only if the 1419 or 1275 contains the dual address adapter. If the single address adapter is used, this operand must be omitted.

BUFFERS= {25|n}: This operand is included to specify the number of buffers in the document buffer area. The limits for n are 12 and 254. 25 is assumed if this operand is omitted.

DEVADDR=SYSnnn: This operand is required and specifies the symbolic unit to be associated with the file. The symbolic unit represents an actual I/O device address used in the ASSGN job control statement to assign the actual I/O device address to the file.

ERROPT=name: This operand may be included only if the CHECK macro is used. The name parameter specifies the name of the routine that the CHECK

macro branches to if any error condition is posted in byte 0, bits 2 to 4 (and bit 5, if no control address is specified in the CHECK macro) of the buffer status indicators. It is your responsibility to exit from this routine (see the "CHECK Macro".)

EXTADDR=name: This operand specifies the name of your stacker selection routine to which control is given when an external interrupt is encountered while reading and sorting the documents internally. This operand may be omitted only when you specify SORTMDE=OFF.

IOAREA1=name: This operand is required and specifies the name of the document buffer area that will be used by the file. Figure 4-3 shows the format of the document buffer area.

IOREG= {(2)|(r)}: This operand specifies the general-purpose register (one of 2 to 12) that the IOCS routines and your routines use to indicate which individual document buffer is available for processing. IOCS puts the address of the current document buffer in the specified register each time a GET or READ is issued. Register 2 is assumed if this operand is omitted.

The same register may be specified in the IOREG entry for two or more files in the same program, if

desired. In this case, your program may need to store the address supplied by IOCS for each record.

MODNAME=name: This operand specifies the name of the logic module generated by MRMOD. If the operand is omitted, IOCS generates the standard system module name.

RECSIZE= {80|n} **:** This operand specifies the actual length of the data portion of the buffer. The record size specified must be the size of the largest record processed. If this operand is omitted, a record size of 80 is assumed. The sum of the ADDAREA and RECSIZE specifications must not exceed 250.

SECADDR=SYSnnn: This operand specifies the symbolic unit to be associated with the secondary control unit address if the 1419 or 1275 with the dual address adapter and LITE macro are utilized. The operand should be omitted if the pocket LITE macro is not being used.

SEPASMB=YES: Include this operand only if the DTFMR is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and defines the filename as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

SORTMDE= {ON|OFF} **:** This operand specifies the method of sorting done on the 1419. SORTMDE=ON indicates that the program sort mode is being used. SORTMDE=OFF indicates that sorting is under control of the magnetic character reader. If the operand is omitted, the program sort mode is assumed.

MRMOD Macro
The first card contains MRMOD in the operation field and may contain a module name in the name field. If a module name is omitted, the following standard module name is generated by IOCS:

$$IJU \begin{Bmatrix} S \\ D \end{Bmatrix} ZZZZ$$

(S = single address adapter, and D = dual address adapter).

The operands you can specify for MRMOD are discussed below.

ADDRESS= {SINGLE|DUAL} **:** Required only if the dual address adapter is used for the 1419 or 1275. If the operand is omitted, the single address adapter is assumed by the assembler.

BUFFERS=nnn: A numeric value (nnn) equal to the corresponding value specified in the DTFMR macro.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and the module name to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

DTFMT Macro

A DTFMT macro is included for each EBCDIC or ASCII magnetic tape input or output file that is to be processed.

Input	Output	Work			
Applies to					►
x	x	x	M	BLKSIZE=nnnnn	Length of one I/O area in bytes (maximum = 32,767).
x	x	x	M	DEVADDR=SYSxxx	Symbolic unit for tape drive used for this file.
x		x	M	EOFADDR=xxxxxxxx	Name of your end-of-file routine.
x	x	x	M	FILABL=xxxx	(NO, STD, or NSTD). If NSTD specified, include LABADDR. If omitted, NO is assumed.
x	x		M	IOAREA1=xxxxxxxx	Name of first I/O area.
x	x		O	ASCII=YES	ASCII file processing is required.
x	x		O	BUFOFF=nn	Length of block prefix if ASCII=YES.
x			O	CKPTREC=YES	Checkpoint records are interspersed with input data records. IOCS bypasses checkpoint records.
x	x	x	O	ERREXT=YES	Additional errors and ERET are desired.
x	x	x	O	ERROPT=xxxxxxxx	(IGNORE, SKIP, or name of error routine). Prevents job termination on error records.
x	x	x	O	HDRINFO=YES	Print header label information if FILABL=STD.
x	x		O	IOAREA2=xxxxxxxx	If two I/O areas are used, the name of the second area.
x	x		O	IOREG=(nn)	Register number. Use only if GET or PUT does not specify a work area or if two I/O areas are used. Omit WORKA. General registers 2-12, written in parentheses.
x	x		O	LABADDR=xxxxxxxx	Name of your label routine if FILABL=NSTD, or if FILABL=STD and user-standard labels are processed.
x			O	LENCHK=YES	Length check of physical records if ASCII=YES and BUFOFF=4.
x	x	x	O	MODNAME=xxxxxxxx	Name of MTMOD logic module for this DTF. If omitted, IOCS generates standard name.
		x	O	NOTEPNT=xxxxxx	(YES or POINTS). YES if NOTE, POINTW, POINTR, or POINTS macro used. POINTS if only POINTS macro used.
x	x	x	O	RDONLY=YES	Generate read-only module. Requires a module save area for each task using the module.
x		x	O	READ=xxxxxxx	(FORWARD or BACK). If omitted, FORWARD assumed.
x	x	x	O	RECFORM=xxxxxx	(FIXUNB, FIXBLK, VARUNB, VARBLK, SPNUNB, SPNBLK, or UNDEF). For work files use FIXUNB or UNDEF. If omitted, FIXUNB is assumed.
x	x		O	RECSIZE=nnnn	If RECFORM=FIXBLK, number of characters in record. If RECFORM=UNDEF, general registers 2-12, written in parentheses. Not required for other records.
x	x	x	O	REWIND=xxxxxx	(UNLOAD or NORWD). Unload on CLOSE or end-of-volume, or prevent rewinding. If omitted, rewind only.
x	x	x	O	SEPASMB=YES	DTFMT is to be assembled separately.
	x		O	TPMARK={YES|NO}	Causes IOCS to write or to omit a tapemark ahead of data records if FILABL=NSTD or NO is specified.
x	x	x	O	TYPEFLE=xxxxxx	(INPUT, OUTPUT, or WORK). If omitted, INPUT is assumed.
	x		O	VARBLD=(nn)	Register number, if RECFORM=VARBLK and records are build in the output area. General registers 2-12, written in parentheses.
x			O	WLRERR=xxxxxxxx	Name of wrong-length-record routine.
x	x		O	WORKA=YES	GET or PUT specifies work area. Omit IOREG.

M=Mandatory
O=Optional

Figure 3-32. DTFMT Macro Operands.

ASCII=YES: This operand specifies that process-ing of ASCII tapes is required. If this operand is omitted, EBCDIC processing is assumed. ASCII=YES is not permitted for work files.

BLKSIZE=nnnnn: Enter the length of the I/O area. If the record format is variable or undefined, enter the length of the largest block of records. If a READ or WRITE macro specifies a length greater than n for work files, the record to be read or written will be truncated to fit in the I/O area. The maximum block size is 32,767 bytes. The minimum size of a physical tape record (gap to gap) is 12 bytes. A record of eleven bytes or less is treated as noise.

For output processing of variable records, the minimum physical record length is 18 bytes. If less than 18 bytes are specified for variable blocked or variable unblocked records, BLKSIZE=18 is assumed.

For output processing of spanned records, the minimum physical record length is 18 bytes. If SPNBLK or SPNUNB and TYPEFLE=OUTPUT are speci-fied in the DTFMT and the BLKSIZE is invalid or less than 18 bytes, an MNOTE is generated and BLKSIZE=18 is assumed.

For ASCII tapes, the BLKSIZE includes the length of any block prefix or padding characters present. If ASCII=YES and BLKSIZE is less than 18 bytes (for fixed-length records only) or greater than 2048 bytes, an MNOTE is generated because this length violates the limits specified by American National Standards Institute, Inc.

BUFOFF= {0|n}: This operand indicates the length of the block prefix. Enter the length of the block prefix if processing of the block prefix is re-quired. This operand can only be included when ASCII=YES is specified. The contents of this field are not passed on to you.

n can have the following values:

Value	Condition
0-99	If TYPEFLE=INPUT
0	IF TYPEFLE=OUTPUT
4	If TYPEFLE=OUTPUT and RECFORM=VARUNB or VARBLK. In this case, the program automatically inserts the physical record length in the block prefix.

CKPTREC=YES: This operand is necessary if an input tape has checkpoint records interspersed among the data records. IOCS bypasses any check-point records encountered. This operand must not be included when ASCII=YES.

DEVADDR= {SYSRDR|SYSIPT|SYSPCH| SYSnnn|SYSLST}:

This operand specifies the symbolic unit to be asso-ciated with the file. An ASSGN job control statement assigns an actual channel and unit number to the unit. The ASSGN job control statement contains the same symbolic name as DEVADDR. When processing ASCII tapes, you must specify a programmer logical unit (SYSnnn).

EOFADDR=name: This operand specifies the name of your end-of-file routine. IOCS automatical-ly branches to this routine on an end-of-file condi-tion. This entry must be specified for input and work files.

In your routine, you can perform any operations required for the end of file (generally you issue the CLOSE macro for the file). IOCS detects end-of-file conditions in magnetic tape input by reading a tape-mark and EOF when standard labels are specified. If standard labels are not specified, IOCS assumes an end-of-file condition when the tapemark is read, if the unit is assigned to SYSRDR or SYSIPT when a /* is read. You must determine, in your routine, that this actually is the end of the file.

ERREXT=YES: This operand enables your ERROPT or WLRERR routine to return to MTMOD by means of the ERET (error return) macro. It also ena-bles nonrecoverable I/O errors occurring before data transfer takes place to be indicated to your program. To take full advantage of this option, the ERROPT=name operand must be specified.

ERROPT= {IGNORE|SKIP|name}: This ope-rand specifies functions to be performed when an error block is encountered. Either IGNORE, SKIP, or the symbolic name of an error routine can be speci-fied. The functions of these specifications are:

IGNORE

The error condition is completely ignored, and the records are made available for processing. When reading spanned records, the entire spanned record or a block of spanned records is returned to the user rather than just the one physical record in which the error occurred.

On output, the error is ignored and the physi-cal record containing the error is treated as a valid record. The remainder, if any, of the spanned record segments are written, if possi-ble.

SKIP

No records in the error block are made avail-able for processing. The next block is read from tape, and processing continues with the first record of that block. The error block is included in the block count. When reading

spanned records, the entire spanned record or a block of spanned records is skipped rather than just one physical record.

On output, the error is ignored and the physical record containing the error is treated as a valid record. The remainder, if any, of the spanned record segments are written.

name

IOCS branches to your error routine named by this parameter regardless of whether ERREXT=YES is specified. In this routine, you can process or make note of the error condition as desired.

The ERROPT operand applies to wrong-length records if the WLRERR operand is not included. If both ERROPT and WLRERR are omitted and wrong-length records occur, IOCS assumes the IGNORE option.

Note: For ASCII tapes, the pointer to the block in error indicates the first logical record following the block prefix.

FILABL= {NO|STD|NSTD}: This operand specifies what type of labels are to be processed. STD indicates standard labels, NO indicates no labels, and NSTD indicates nonstandard labels. You must furnish a routine to check or create the nonstandard labels by using your own I/O area and an EXCP macro to read or write the labels. The entry point of this routine is the operand of LABADDR.

The specification FILABL=NSTD is not permitted for ASCII files (that is, when ASCII=YES). Labels and tape data are assumed to be in the same mode.

HDRINFO=YES: This operand, if specified with FILABL=STD, causes IOCS to print standard header label information (fields 3-10) on SYSLOG each time a file with standard labels is opened. It also prints the filename, logical unit, and device address each time an end-of-volume condition is detected. Both FILABL=STD and HDRINFO=YES must be specified for header label information to be printed.

IOAREA1=name: This operand specifies the name of the I/O area. When variable-length records are processed, the size of the I/O area must include four bytes for the block size. This operand does not apply to work files.

IOAREA2=name: This operand specifies the name of a second I/O area. When variable-length records are processed, the size of the I/O area must include four bytes for the blocksize. This operand does not apply to work files.

IOREG=(r): This operand specifies the register in which IOCS places the address of the logical record that is available for processing if:

- two input or output areas are used.
- blocked input or output records are processed in the I/O area.
- variable unblocked records are read.
- undefined records are read backwards.
- neither BUFOFF=0 nor WORKA=YES is specified for ASCII files.

For output files, IOCS places, in the specified register, the address of the area where you can build a record. Any of registers 2 to 12 may be specified.

This operand cannot be used if WORKA=YES.

LABADDR=name: Enter the symbolic name of your routine to process user-standard or nonstandard labels.

For ASCII tapes, this operand may be used only for writing and checking user standard labels that conform to American National Standards Institute, Inc., standards. You must process these labels in EBCDIC. Nonstandard user labels are not permitted.

LENCHK=YES: This operand applies only to ASCII tape input if BUFOFF=4 and RECFORM=VARUNB or VARBLK. It must be included if the block length (specified in the block prefix) is to be checked against the physical record length. If the two lengths do not match, the action taken is the same as described under the WLRERR operand, but the WLR bit (byte 5, bit 1) in the DTF is not set.

MODNAME=name: This operand specifies the name of the logic module used with the DTF table to process the file. If the logic module is assembled with the program, this operand must specify the same name as the MODNAME operand of the MTMOD macro. If this operand is omitted, standard names are generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called. For example, if one DTF specifies READ=FORWARD and another specifies READ=BACK, only one logic module capable of handling both functions is called.

NOTEPNT= {POINTS|YES}: If the parameter YES is specified, the NOTE, POINTW, POINTR, or POINTS macros can be issued for a tape work file. If POINTS is specified, only POINTS macros can be issued for tape work files. The NOTEPNT operand must not be specified for ASCII tape files because ASCII work files are not supported by DOS/VSE.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module.

READ= {FORWARD|BACK}: This operand specifies the direction in which the tape is read. If READ=BACK is specified and a wrong-length record smaller than the I/O area is encountered, the record is read into the I/O area right-justified.

RECFORM= {FIXUNB|FIXBLK|VARUNB|
     VARBLK|SPNBLK|SPNUNB|
     UNDEF}:
This operand specifies the type of EBCDIC or ASCII records in the input or output file. Enter one of the following parameters:

FIXUNB	For fixed-length unblocked records
FIXBLK	For fixed-length blocked records
VARUNB	For variable-length unblocked records
VARBLK	For variable-length blocked records
SPNBLK	For spanned variable-length blocked records (EBCDIC only)
SPNUNB	For spanned variable-length unblocked records (EBCDIC only)
UNDEF	For undefined records

Work files may use only FIXUNB or UNDEF.

RECSIZE= {n|(r)} For fixed-length blocked records, RECSIZE is required. It specifies the number of characters in each record.

When processing spanned records, you must specify RECSIZE=(r) where r is a register that contains the length of each record.

For undefined records, this entry is required for output files but is optional for input files. It specifies a general register (any of 2 to 12) that contains the length of the record. On output, you must load the length of each record into the register before you issue a PUT macro. Spanned-record output requires a minimum record length of 18 bytes. A physical record less than 18 bytes is padded with binary zeros to complete the 18-byte requirement. This applies to both blocked and unblocked records. If specified for input, IOCS provides the length of the record transferred to virtual storage.

REWIND= {UNLOAD|NORWD}: If this specification is not included, tapes are automatically rewound to load point, but not unloaded, on an OPEN or OPENR or a CLOSE or CLOSER macro or on an end-of-volume condition. If other operations are desired for a tape input or output file, specify:

UNLOAD
     to rewind the tape on an OPEN (or OPENR) and
     to rewind and unload on a CLOSE (or CLOSER)
     or on an end-of-volume condition.

NORWD
     to prevent rewinding the tape at any time.
     This option positions the read/write head between the two tapemarks that indicate the end-of-file condition.

SEPASMB=YES: Include this operand only if the DTFMT is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TPMARK= {YES|NO}: A tapemark is normally written for an output file if nonstandard labels are specified (FILABL=NSTD). If no tapemark is desired, TPMARK=NO should be specified. If TPMARK=NO is specified together with FILABL=STD, the former specification is ignored. If FILABL=NO is specified or the FILABL operand is omitted, TPMARK=YES must be specified for IOCS to write a tapemark ahead of the first data record.

TYPEFLE= {INPUT|OUTPUT|WORK}: Use this operand to indicate whether the file is used for input or output. If INPUT is specified, the GET macro is used. If OUTPUT is specified, the PUT macro is used. If WORK is specified, the READ/WRITE, NOTE/POINTx, and CHECK macros are used.

The specification of WORK in this operand is not permitted for ASCII files.

VARBLD=(r): This entry is required whenever variable-length blocked records are built directly in the output area (no work area is specified). It specifies the number (r) of a general-purpose register (any of 2 to 12) that always contains the length of the available space remaining in the output area.

IOCS calculates the space still available in the output area, and supplies it to you in the VARBLD register after the PUT macro is issued for a variable-length record. You can then compare the length of the next variable-length record with the available space to determine whether the record will fit in the remaining area. This check must be made before the record is built. If the record does not fit, issue a TRUNC macro to transfer the completed block of records to the tape. The current record is then built as the first record of the next block.

WLRERR=name: This operand applies only to tape input files. It specifies the name of your routine to receive control if a wrong-length record is read. If

the WLRERR entry is omitted but a wrong-length record is detected by IOCS, one of the following conditions results:

- If the ERROPT entry is included for this file, the wrong-length record is treated as an error block, and handled according to your specifications for an error (IGNORE, SKIP, or name of error routine).

- If the ERROPT entry is not included, IOCS assumes the IGNORE option of ERROPT.

WORKA=YES: If I/O records are processed in work areas instead of in the I/O areas, specify this operand. You must set up the work areas in virtual storage. The symbolic address of the work area, or a general-purpose register containing the address, must be specified in each GET or PUT. Omit IOREG if this operand is included. WORKA=YES is required for spanned record processing.

MTMOD Macro

Listed here are the operands you can specify for MTMOD. The first card contains MTMOD in the operation field and may contain a module name in the name field.

ASCII=YES: Include the operand if processing ASCII input or output files is required. This entry is not permitted for work files. If omitted, EBCDIC file processing is assumed.

CKPTREC=YES: Include this operand if input tapes processed by the module contain checkpoint records interspersed among the data records. The module also processes tapes that do not have checkpoint records; that is, those whose DTFs do not specify CKPTREC=YES.
This entry is not needed for work files, and is not valid for ASCII files.

ERREXT=YES: Include this operand if additional I/O errors are to be indicated and/or the ERET macro is used with this DTF and module. ERROPT=YES should be specified in this module for work files, but is not needed for input or output files.

ERROPT=YES: Include this operand if the module is to handle any of the error options for an error block. Code is generated to handle any of the three options (IGNORE, SKIP, or name). The module also processes any files in which the ERROPT operand is not specified in the DTF. This entry is needed for work files, but it is not needed for input or output files.

NOTEPNT={YES|POINTS}: This operand applies only to work files (EBCDIC only).

Include this operand if NOTE/POINTx logic is used with the module. If YES is specified, the module processes any NOTE, POINTR, POINTW, or POINTS macro. If POINTS is specified, only the POINTS macro is processed.

Modules specifying either one of the two options also process work files for which the NOTE/POINTx operand is not specified in the DTF. Modules specifying YES also process work files specifying only POINTS.

RDONLY=YES: This operand causes a read-only module to be generated. Whenever this operand is specified, any DTF used with the module must have the same operand.

Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each task should have its own uniquely defined save area and each time an imperative macro (except an OPEN, OPENR, or LBRET) is issued, register 13 must contain the address of the save area associated with the task. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If the operand is omitted, the module generated is not reenterable and no save area is required.

READ={FORWARD|BACK}: This operand generates a module that reads tape files forward or backward. If forward is specified, only code to read tape forward is generated. Any DTF used with the module must not specify BACK in the READ parameter statement.

If the parameter is BACK, code to read tape both forward and backward is generated, and any DTF used with the module may specify either FORWARD or BACK as its READ parameter. READ=BACK does not handle multi-volume files.

This entry is not needed for work files.

RECFORM={FIXUNB|FIXBLK|VARUNB|VARBLK|SPNBLK|SPNUNB|UNDEF}:
This operand generates an input/output module that processes either EBCDIC or ASCII fixed-length, variable-length, or undefined records.

If FIXUNB or FIXBLK is specified, a module is generated that allows processing of *both* fixed-length blocked and fixed-length unblocked records. Similarly, if VARUNB/SPNUNB or VARBLK/SPNBLK is specified, a module is generated that allows processing of both types of variable and spanned records. ASCII files are not permitted in spanned record format.

If UNDEF is specified, a module for processing undefined record types is generated. Any DTF used with the module must specify the same record format type as the module. For example, if the module specifies RECFORM=FIXUNB, either RECFORM=FIXUNB or RECFORM=FIXBLK may be specified in the DTF.

This operand is not needed for work files.

If this operand is omitted, the module generated will allow processing of both fixed-length blocked and fixed-length unblocked records.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and the module to be defined as an ENTRY

point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

TYPEFLE= {OUTPUT|INPUT|WORK}: This operand generates a module that processes either GET/PUT macros or READ/WRITE, NOTE/POINTx and CHECK macros for work files (EBCDIC only). If the parameter of the operand specifies WORK, code to process work files is generated. Otherwise, a module to handle both input and output files is assumed. Only DTFs for work files may be used with work file modules. Only DTFs for input or output files may be used with an input/output module.

Note: INPUT and OUTPUT have the same table format and logic modules.

WORKA=YES: This operand is to be included if records are to be processed in work areas instead of I/O areas for the GET/PUT macros. This operand must be included if spanned records are processed. The module also processes files that do not use a work area. This entry is not needed for work files.

Standard MTMOD Names

Each name begins with a 3-character prefix (IJF) and continues with a 5-character field containing the options permitted in module generation. In MTMOD there are two module classes: the module class for handling GET/PUT functions and the module class for handling READ/WRITE, NOTE/POINTx, and CHECK functions (work files). Modules handling fixed-length (F and X) and undefined (U and N) records are mutually exclusive of each other and of all forms of the module that process variable-length records (V, R, and S).

Name list for GET/PUT type modules:

MTMOD name = IJFabcde

a = F RECFORM=FIXUNB (or FIXBLK) (EBCDIC mode)
  = X RECFORM=FIXUNB (or FIXBLK) (ASCII mode)
  = V RECFORM=VARUNB (or VARBLK) (EBCDIC mode)
  = R RECFORM=VARUNB (or VARBLK) (ASCII mode)
  = S RECFORM=SPNUNB (or SPNBLK) (spanned records)
  = U RECFORM=UNDEF (EBCDIC mode)
  = N RECFORM=UNDEF (ASCII mode)

b = B READ=BACK
  = Z READ=FORWARD, or if READ is not specified

c = C CKPTREC=YES
  = Z CKPTREC=YES is not specified

d = W WORKA=YES
  = Z WORKA=YES is not specified

e = M ERREXT=YES and RDONLY=YES
  = N ERREXT=YES
  = Y RDONLY=YES
  = Z ERREXT and RDONLY not specified

Name list for work file type modules (TYPEFLE=WORK):

MTMOD name = IJFabcde

a = W

b = E ERROPT=YES
  = Z ERROPT is not specified

c = N NOTEPNT=YES
  = S NOTEPNT=POINTS
  = Z NOTEPNT is not specified

d = Z always

e = M ERREXT=YES and RDONLY=YES
  = N ERREXT=YES
  = Y RDONLY=YES
  = Z ERREXT and RDONLY not specified

Subset/Superset MTMOD Names

The charts in Figure 3-33 illustrate the subsetting and supersetting allowed for MTMOD names. Four of the GET/PUT parameters allow subsetting. For example, the module name IJFFBCWZ is a superset of IJFFBZWZ specifying fixed-length records.

```
For GET/PUT Type Modules:
        *  +  +  +  +
  I J F  F  B  C  W  M
         N  Z  Z  Z  Y
         R           +
         U           N
         X           Z
         +
         S
         V
For Workfile Type Modules:
           +  +     +
  I J F W  E  N  Z  M
           Z  S     Y
              Z     +
                    N
                    Z
```

| + Subsetting/supersetting permitted. |
| * No subsetting/supersetting permitted. |

Figure 3-33. Subsetting and supersetting of MTMOD names.

DTFOR Macro

DTFOR is used to define an input file to be processed on a 1287 optical reader or 1288 optical page reader. The macro is not used for the 3881 optical mark reader; for the 3881, use the DTFCD macro.

Applies to					
1287T	1287D	1288			
x	x	x	M	COREXIT=xxxxxxx	Name of your correction routine
x	x	x	M	DEVADDR=SYSnnn	Symbolic unit assigned to the optical reader
x	x	x	M	EOFADDR=xxxxxxx	Name of your end-of-file routine
x	x	x	M	IOAREA1=xxxxxxx	Name of first input area
x			O	BLKFAC=nn	If RECFORM=UNDEF in journal tape mode
x	x	x	O	BLKSIZE=nn	Length of I/O area(s). If omitted, 38 is assumed.
x	x	x	O	CONTROL=YES	If CNTRL macro is to be used for this file
x	x	x	O	DEVICE=xxxxx	(1287D or 1287T). For 1288, specify 1287D. If omitted, 1287D is assumed.
x	x		O	HEADER=YES	If a header record is to be read from the optical reader keyboard by OPEN, OPENR
	x	x	O	HPRMTY=YES	If hopper empty condition is to be returned.
x			O	IOAREA2=xxxxxxx	If two input areas are used, name of second input area.
x			O	IOREG=(nn)	Register number if 2 input areas or UNDEF records are to be used. If omitted, register 2 is assumed. General registers 2-12, written in parentheses.
x	x	x	O	MODNAME=xxxxxxx	Name of DTF's logic module. If omitted, IOCS generates a standard name.
x	x	x	O	RECFORM=xxxxx	(FIXBLK, FIXUNB, or UNDEF). If omitted, FIXUNB is assumed.
x	x	x	O	RECSIZE=(nn)	Register number containing record size, if RECFORM=UNDEF. If omitted, register 3 is assumed.
x	x	x	O	SEPASMB=YES	If the DTFOR is to be assembled separately.
x			O	WORKA=YES	If records are to be processed in a work area. Omit IOREG.

M=Mandatory
O=Optional

Figure 3-34. DTFOR macro options.

BLKFAC=n: Undefined journal tape records are processed with greater throughput speeds when this operand is included. This is accomplished by reading groups of lines as blocked records. When undefined records are processed, BLKFAC specifies the blocking factor (n) that determines the number of lines read (through CCW chaining) as a block of data by one physical read. Deblocking is accomplished automatically by IOCS when the GET macro is used. The BLKFAC parameter is not used with RECFORM=FIXBLK, because the blocking factor is determined from the BLKSIZE and RECSIZE parameters. If the operand is included for FIXBLK, FIXUNB, or document processing, the operand is noted (in an MNOTE) and ignored.

BLKSIZE={38|n}: This operand indicates the size of the input area specified by IOAREA1. For journal tape processing, BLKSIZE specifies the maximum number of characters that can be transferred to the area at any one time.

When undefined journal tape records are read, the area must be large enough to accommodate the longest record to be read if the BLKFAC parameter is not specified. If the BLKFAC parameter is specified, the BLKSIZE value must be determined by multiplying the maximum length that must be accommodated for an undefined record by the blocking factor desired. A BLKSIZE value smaller than this results in truncated data.

If two input areas are used for journal tape processing (IOAREA1 and IOAREA2), the size specified in this entry is the size of each I/O area.

CONTROL=YES: This entry must be included if a CNTRL macro is issued for a file. A CNTRL macro issues orders to the optical reader to perform nondata operations such as line marking, stacker selecting, document incrementing, etc.

COREXIT=name: COREXIT provides an exit to your error correction routine for the 1287 or 1288. After a GET, WAITF, or CNTRL macro is executed (to increment or eject and/or stacker select a document), an error condition causes an error correction routine to be entered with an error indication provided in filename+80. The byte at filename+80 contains the following codes indicating the conditions that occurred during the last line or field read. The byte should also be tested after issuing the optical reader macros DSPLY, RESCN, RDLNE, CNTRL READKB, and CNTRL MARK. More than one error condition may be present.

Code		Meaning
Dec	Hex	
1	X'01'	A data check has occurred. Five read attempts for journal tape processing or three read attempts for document processing were made.
2	X'02'	The operator corrected one or more characters from the keyboard (1287T) or a hopper empty condition (see HPRMTY=YES operand) has occurred (1287D).
4	X'04'	A wrong-length record condition has occurred (for journal tapes, five read attempts were made; for documents, three read attempts were made). Not applicable for undefined records.
8	X'08'	An equipment check resulted in an incomplete read (ten read attempts were made for journal tapes or three for documents).
		If an equipment check occurs on the first character in the record, when processing undefined journal tape records, the RECSIZE register contains zero, and the IOREG (if used) points to the rightmost position of the record in the I/O area. You should test the RECSIZE register before moving records from the work area or the I/O area.
16	X'10'	A nonrecoverable error occurred.
32	X'20'	For the 1288, reading in unformatted mode, the end-of-page (EOP) condition has been detected. Normally, on an EOP indication, the problem program ejects and stacker selects the document.
		After issuing one of the macros CNTRL ESD, CNTRL SSD, CNTRL EJD in your COREXIT routine, a late stacker selection condition occurred.
		For the 1287, a stacker select was given after the allotted elapsed time and the document was put in the reject pocket.
64	X'40'	The 1287D scanner was unable to locate the reference mark (for journal tapes, ten read attempts were made; for documents, three read attempts were made).

The byte filename+80 can be interrogated to determine the reason for entering the error correction routine. Choice of action in your error correction routine is determined by the particular application.

If you issue I/O macros to any device other than the 1287 and/or 1288 in the COREXIT routine, you must save registers 0, 1, 14, and 15 upon entering the routine, and restore these registers before exiting. Furthermore, if I/O macros (other than the GET, WAITF, and/or READ, which cannot be used in COREXIT) are issued to the 1287 and/or 1288 in this routine, you must also save and later restore registers 14 and 15 before exiting. All exits from COREXIT should be to the address specified in register 14. This provides a return to the point from which the branch to COREXIT occurred. If the command chain bit is on in the READ CCW for which the error occurred, IOCS completes the chain upon return from the COREXIT routine.

Note: Do not issue a GET, READ, OPEN, or WAITF macro to the 1287 or 1288 in the error correction routine. Do not process records in the error correction routine. The record that caused the exit to the error routine is available for processing upon return to the mainline program. Any processing included in the error routine would be duplicated after return to the mainline program.

When processing journal tapes, a nonrecovery error (torn tape, tape jam, etc.) normally requires that the tape be completely reprocessed. In this case, your routine **must not** branch to the address in register 14 from the COREXIT routine or a program loop will occur. Following an unrecoverable error:

- the optical reader file must be closed.
- the condition causing the nonrecovery must be cleared.
- the file must be reopened before processing can continue.

If a nonrecoverable error occurs while processing documents (indicating that a jam occurred during a document incrementation operation, or a scanner control failure has occurred, or an end-of-page condition, etc.), the document should be removed either manually or by nonprocess runout. In such cases, your program should branch to read the next document.

If the 1287 or 1288 scanner is unable to locate the document reference mark, the document cannot be processed. In this case, the document must be ejected and stacker selected before attempting to read the following document or a program loop will result.

Whenever a nonrecoverable error occurs, your COREXIT routine **must not** branch to the address in register 14 to return to IOCS. Instead, the routine should ignore any output resulting from the document.

Eight binary error counters are used to accumulate totals of certain 1287 and 1288 error conditions. Each of these counters occupies four bytes, starting at filename+48. Filename is the name specified in the DTF header entry. The error counters are:

Counter and Address	Contents
1 filename+48	Equipment check (see Note, below).
2 filename+52	Equipment check uncorrectable after ten read attempts for journal tapes or three read attempts for documents (see Note, below).
3 filename+56	Wrong-length records (not applicable for undefined records).
4 filename+60	Wrong-length records uncorrectable after five read attempts for journal tapes or three read attempts for documents (not applicable for undefined records).
5 filename+64	Keyboard corrections (journal tape only).
6 filename+68	Journal tape lines (including retried lines) or document fields (including retried fields) in which data checks are present.
7 filename+72	Lines marked (journal tape only).
8 filename+76	Count of total lines read from journal tape or the number of CCW chains executing during document processing.

Note: Counters 1 and 2 apply to equipment checks that result from incomplete reads or from the inability of the 1287 or 1288 scanner to locate a reference mark (when processing documents only).

All the previous counters contain binary zeros at the start of each job step. You may list the contents of these counters for analysis at end of file, or at end of job, or you may ignore the counters. The binary contents of the counters should be converted to a printable format.

DEVADDR=SYSnnn: This operand specifies the logical unit (SYSnnn) to be associated with the file. The logical unit represents an actual I/O device address used in the ASSGN job control statement to assign the actual I/O device address to this file.

DEVICE={1287D|1287T}: This operand specifies the I/O device associated with this file. 1287D specifies a 1287 or 1288 document file. 1287T specifies a 1287 journal tape file.

From this specification, IOCS sets up the device-dependent routines for this file. For document processing you must code the CCWs.

If this operand is omitted, 1287D is assumed.

EOFADDR=name: This operand specifies the name of your end-of-file routine. IOCS automatically branches to this routine on an end-of-file condition.

When reading data from documents, you can recognize an end-of-file condition by pressing the end-of-file key on the console when the hopper is empty. When processing journal tapes on a 1287, you can detect an end-of-file by pressing the end-of-file key after the end of the tape is sensed.

When IOCS detects an end-of-file condition, it branches to your routine specified by EOFADDR. You must determine whether the current roll is the last roll to be processed when handling journal tapes. Regardless of the situation, the tape file must be closed for each roll within your EOF routine. If the current roll is not the last, OPEN (or OPENR) must be issued. The OPEN (or OPENR) macro allows header (identifying) information to be entered at the reader keyboard and read by the processor when using logical IOCS.

The same procedure can be used for 1287 processing of multiple journal tape rolls, as well as the method described under "OPEN (or OPENR) Macro" in the section "Imperative Macros".

HEADER=YES: This operand cannot be used for 1288 files. This operand is required if the operator is to key in header (identifying) information from the 1287 keyboard. The OPEN (or OPENR) routine reads the header information only when this entry is present. If the entry is not included, OPEN (or OPENR) assumes no header information is to be read. The header record size can be as large as the BLKSIZE entry and is read into the high-order positions of IOAREA1.

HPRMTY=YES: This operand is included if you want to be informed of the hopper empty condition. This condition occurs when a READ is issued and no document is present, and is recognized at WAITF time. When a hopper empty condition is detected, your COREXIT routine is entered with X'02' stored in filename+80.

This operand should be used when processing documents in the time-dependent mode of operation, which allows complete overlapping of processing with reading. See the appropriate IBM 1287 device manuals for processing details. With this method of processing, specifying HPRMTY=YES allows you to check for a hopper empty condition in your COREXIT routine. You can then select into the proper hopper the previously ejected document before return from COREXIT (via register 14).

IOAREA1=name: This operand is included to specify the name of the input area used by the file. When opening a file and before each journal tape input operation to this area, the designated area is set to binary zeros and the input routines then transfer records to this area. For document processing, the area is cleared only when the file is opened.

IOAREA2=name: A second input area can be allotted only for a journal tape file. This permits an overlap of data transfer and processing operations. The specified second I/O area is set to binary zeros before each input operation to this area occurs.

IOREG= {(2)|(r)}: This operand specifies a general-purpose register (any one of 2 to 12) that the input routines use to indicate the beginning of records for a journal tape file. The same register may be specified in the IOREG operand for two or more files in the same program, if desired. In this case, your program may need to store the address supplied by IOCS for each record. Whenever this entry is included for a file, the DTFOR entry WORKA must be omitted, and the GET macro must not specify a work area.

A read by an optical reader is accomplished by a backward scan. This places the rightmost character in the record in the rightmost position in the I/O area and subsequent characters in sequence from right to left. The register defined by IOREG indicates the leftmost position of the record.

MODNAME=name: This operand may be used to specify the name of the logic module used with the DTF table to process the file. If the logic module (ORMOD) is assembled with the program, the MODNAME parameter in this DTF must specify the same name as the ORMOD macro.

If this entry is omitted, standard names are generated for calling the logic module. If two different DTF macros call for different functions that can be handled by a single module, only one standard-named module is called.

RECFORM= {FIXUNB|FIXBLK|UNDEF}:
This operand specifies the type of records in an optical reader file. One of the following may be specified:

FIXUNB   For fixed-length unblocked records.
FIXBLK   For fixed-blocked records in journal tape mode.
UNDEF   For undefined records.

RECSIZE= n| {(3)|(r)}: For fixed-length un-

blocked records, this operand should be omitted and no register is assumed.

For fixed-length blocked records (journal tape mode), this operand must be included to specify the number, n, of characters in an individual record. The input routines use this number to deblock records, and to check the length of input records. If this operand is omitted, an MNOTE is flagged in the macro assembly and fixed-length unblocked records are assumed.

For undefined journal tape records, this entry specifies the number (r) of the general-purpose register in which IOCS provides the length of each input record. For undefined document records, RECSIZE contains only the length of the last field of a document read by the CCW chain that you supply. Any one of registers 2 through 12 may be specified, but if the operand is omitted, register 3 is assumed.

Note: When processing undefined records in document mode, you gain complete usage of the register normally used in the RECSIZE operand. You can do this by ensuring that the suppress-length-indication (SLI) flag is always on when processing undefined records.

SEPASMB=YES: Include this operand only if the DTFOR will be assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

WORKA=YES: Input records (journal tape only) can be processed in work areas instead of in the input areas. If this is planned, the operand WORKA=YES must be specified, and you must set up the work area in storage. The symbolic name of the work area, or a general-purpose register containing the address of the work area, must be specified in each GET macro. When GET is issued, IOCS left-justifies the record in the specified work area. Whenever this operand is included for a file, the DTFOR IOREG operand must be omitted.

ORMOD Macro

Listed here are the operands you can specify for ORMOD. The first card contains ORMOD in the operation field and may contain a module name in the name field.

Note: ORMOD is not used for the 3881 Optical Mark Reader. The 3881 uses CDMOD.

BLKFAC=YES: Include this operand if RECFORM=UNDEF and groups of undefined journal tape records are to be processed as blocks of data. (See the DTFOR BLKFAC=n operand.) The DTFOR used with this module must also include RECFORM=UNDEF and BLKFAC=n.

CONTROL=YES: Include this operand if CNTRL macros are to be used with the associated DTFs. The module also processes files that do not use the CNTRL macro.

DEVICE= {1287D|1287T}: This operand must be included to specify the I/O device associated with this file. 1287D specifies a 1287 or 1288 document file. 1287T specifies a 1287 journal tape file.

IOAREA2=YES: Include this operand (journal tape only) if a second I/O area is used. The DTFOR used with this module must also include the IOAREA2 parameter.

RECFORM= {FIXUNB|FIXBLK|UNDEF}: This operand generates a module that processes the specified record format. Any DTF used with the module must have the same operand.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and the module name to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the ORMOD macro is being assembled with the DTF and the problem program and no CATALR card is punched.

WORKA=YES: Include this operand (journal tape only) if records are to be processed in work areas instead of in I/O areas. Any DTF used with the module must have the same operand.

Standard ORMOD Names

Each name begins with a 3-character prefix (IJM) followed by a 5-character field corresponding to the options permitted in the generation of the module.

ORMOD name = IJMabcde

```
a  =  F  RECFORM=FIXUNB
   =  X  RECFORM=FIXBLK
   =  U  RECFORM=UNDEF
   =  D  RECFORM=UNDEF and BLKFAC=YES
b  =  C  CONTROL=YES
   =  Z  CONTROL=YES is not specified
c  =  I  IOAREA2=YES
   =  W  WORKA=YES
   =  B  both are specified
   =  Z  neither is specified
d  =  T  device is in tape mode
   =  D  device is in document mode
e  =  Z  always
```

Subset/Superset ORMOD Names

Figure 3-35 shows the subsetting and supersetting allowed for ORMOD names. One of the parameters allows subsetting. For example, the module IJMFCITZ is a superset of the module IJMFZITZ.

```
          *  +  *  *
    I  J  M  D  C  B  D  Z
             F  Z  I  T
             U     W
             X     Z
```
+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.

Figure 3-35. Subsetting and supersetting of ORMOD names.

DTFPH Macro

When physical IOCS macros (EXCP, WAIT, etc.) are used in a program, DASD, diskette, or tape files with standard labels need to be defined by the DTFPH macro (DTFxx macro for a file handled by physical IOCS). DTFPH must also be used for a checkpoint file on a disk.

M	TYPEFLE=xxxxxx	(INPUT or OUTPUT). Specifies type of file.
O	ASCII=YES	ASCII file processing is required.
O	CISIZE=n	For Fixed Block Architecture DASD, Control Interval size.
O	CCWADDR=xxxxxxxx	If CCB is generated by DTFPH is to be used.
O	DEVICE=xxxx	(TAPE, FBA, 2311, 2314, 3330, 3340, 3350, 3540). If omitted, TAPE is assumed.
O	DEVADDR=SYSxxx	Symbolic unit required only when not provided on an EXTENT statement.
O	HDRINFO=YES	Print header label information.
O	LABADDR=xxxxxxxx	Routine to check or build user standard labels.
O	MOUNTED=xxxxxx	(ALL or SINGLE). Required for DASD files only; for diskette files, specify SINGLE.
O	XTNTXIT=xxxxxxxx	If EXTENT statements are to be processed. DASD only.

M=Mandatory
O=Optional

Figure 3-36. DTFPH macro.

Figure 3-37 shows which of the DTFPH entries can or must be coded to define a checkpoint file on disk.

Operand	Optional	Required
CCWADDR=name	x	
CISIZE=n	x	
DEVADDR=SYSnnn	x	
DEVICE=2311, 2314, 3330, 3340, 3350, 3540, FBA		x
LABADDR=name	x	
MOUNTED=SINGLE		x
TYPEFLE=OUTPUT		x

Figure 3-37. Operands to define a checkpoint file on disk.

ASCII=YES: This operand is required to process ASCII tape files. If this operand is omitted, EBCDIC processing is assumed.

CCWADDR=name: This operand allows you to use the CCB generated within the first 16 bytes of the DTFPH table. CCWADDR specifies the symbolic name of the first CCW used with the CCB generated within the DTFPH macro. This name must be the same as the name specified in the assembler CCW statement that constructs the CCW.

If this operand is omitted, the location counter value of the CCB-CCW table address constant is substituted for the CCW address.

CISIZE=n: This operand specifies the FBA Control Interval size. The value n must be an integral multiple of the FBA logical block size and, if greater than 8K, must be a multiple of 2K. The maximum value is 32768 (32K) except when assigned to SYSLST or SYSPCH, when the maximum is 30720 (30K).

If DEVICE=FBA is specified, and CISIZE is omitted, CISIZE=0 is assumed. Control Interval size may be overridden for an output file at execution time by specifying the CISIZE parameter of the DLBL job control statement. For an input file, the CISIZE value in the format-1 label is used.

DEVICE= {TAPE|FBA|2311|2314|3330|3340| 3350|3540}:
If the file is contained on DASD or diskette, enter the proper identification.

TAPE applies to 8809 and any 2400/3400-series tape unit, and is the only valid entry in this operand for ASCII files.

FBA applies to 3310 and 3370.

For devices supported by DOS/VSE and not included in the above operand specification, specify device codes as listed in Figure 3-38.

DEVADDR=SYSxxx: This operand must specify the symbolic unit (SYSxxx) associated with the file if a symbolic unit is not provided via an EXTENT job control statement. If a symbolic unit is provided, its specification overrides a DEVADDR specification. This specification, or symbolic unit, represents an actual I/O address, and is used in the ASSGN job control statement to assign the actual I/O device address to this file.

If SYSLST or SYSPCH are used as output tape units and alternate tape switching is desired upon detect-

DEVICE = specification	Device in use									
	2311	2314	2319	3330-1,2*	3330-11**	3340, 35MB	3340, 70MB	3350	TAPE	3310, 3370
TAPE									x	
2311	x				x			x		x
2314		x	x		x			x		x
3330				x	x			x		x
3340					x	x	x	x		x
3350					x			x		x
FBA										x

    * Also 3350 in 3330-1 compatibility mode.
    ** Also 3350 in 3330-11 compatibility mode.

Figure 3-38. DEVICE= specifications for DTFPH.

ing a reflective spot, the SEOV macro must be used (see "SEOV Macro"). When processing ASCII tape files, the only valid specification is a programmer logical unit (that is, SYSnnn).

HDRINFO=YES: This operand causes IOCS to print standard header label information (fields 3-10) on SYSLOG each time a file with standard labels is opened. Likewise, the filename, symbolic unit, and device address are printed each time an end-of-volume condition is detected. If HDRINFO=YES is omitted, no header or end-of-volume information is printed.

LABADDR=name: This operand does not apply to diskette input/output units.

You may require one or more DASD or tape labels in addition to the standard file labels. If so, you must include your own routine to check (on input) or build (on output) your label(s). Specify the symbolic name of your routine in this operand. IOCS branches to this routine after the standard label is processed.

LABADDR may be included to specify a routine for your header or trailer labels as follows:

- DASD input or output: header labels only.
- Tape input or output: header and trailer labels.

Thus, if LABADDR is specified, your header labels can be processed for an input/output DASD or tape file, and your trailer labels can be built for a tape output file. Physical IOCS reads input labels and makes them available to you for checking, and writes output labels after they are built. This is similar to the functions performed by logical IOCS.

If physical IOCS macros are used for a tape file, an OPEN must be issued for the new volume. This caus-

es IOCS to check the HDR1 label and provides for your checking of user standard labels, if any.

When physical IOCS macros are used and DTFPH is specified for standard tape label processing, FEOV must not be issued for an input file.

MOUNTED= {ALL|SINGLE} This operand does not apply to diskette input/output units.

This operand must be included to specify how many extents (areas) of the file are available for processing when the file is initially opened. This operand must not be specified for tape.

ALL is specified if all extents are available for processing. When a file is opened, IOCS checks all labels on each disk pack and makes available all extents specified by your control statements. Only one OPEN or OPENR is required for the file. ALL should be specified whenever you plan to process records in a manner similar to the direct access method. In any case, you must supply an LBLTYP statement.

Note: On systems with VSE/Advanced Functions installed, the LBLTYP statement is not required.

After an OPEN or OPENR is performed, you must be aware that the symbolic unit address of the first volume containing the file is in bytes 30 and 31 of the DTFPH table rather than in the CCB. Therefore, place this symbolic address into bytes 6 and 7 of the associated CCB before you issue an EXCP against this CCB in your program.

SINGLE is specified if only the first extent on the first volume is available for processing. SINGLE should be specified when you plan to process records in sequential order. IOCS checks the labels on the first pack and makes the first extent specified by your control cards available for processing. You must keep track of the extents and issue a subse-

quent OPEN or OPENR whenever another extent is required for processing. You will find the information in the DTFPH table helpful in keeping track of the extents. The DTFPH table contains:

Bytes	Contents
0-15	CCB (symbolic unit has been initialized in the CCB).
54-57	Extent upper limits (cchh).
58-59	Seek address. For a disk it must be zero.
60-63	Extent lower limit (cchh).

On each OPEN or OPENR after the first, IOCS makes available the next extent specified by the control cards. When you issue a CLOSE or CLOSER for an output file, the volume on which you are currently writing records is indicated, in the file label, as the last volume for the file.

TYPEFLE= {INPUT|OUTPUT}: This operand must be included to specify the type of file: input or output.

XTNTXIT=name: This operand does not apply to diskette input/output units.

This entry is included if you want to process label extent information. It specifies the symbolic name of your extent routine. The DTFPH operand MOUNTED=ALL must also be specified for the file.

Whenever XTNTXIT is included, IOCS branches to your routine during the initial OPEN for the file. It branches after each specified extent is completely checked and after conflicts, if any, have been resolved.

When your routine receives control, register 1 contains the address of a 14-byte area from which you can retrieve label extent information (in binary form). The layout of this area is shown in Figure 3-39.

Return to IOCS by using the LBRET macro.

Bytes	Contents	
0	Extent type code:	
	00	Next three fields do not indicate any extent.
	01	The extent containing your data record.
	02	Overflow area of an indexed sequential file.
	04	Cylinder index or master index of an indexed sequential file.
	40	User label track area
	80	Shared cylinder indicator
1	Extent sequence number.	
2-5	Lower limit of the extent (cchh).	
6-9	Upper limit of the extent (cchh).	
10-11	Symbolic unit (see Figure 2-1).	
12	Contains zero.	
13	Not used.	

Figure 3-39. Layout of XTNTXIT information area.

DTFPR Macro
DTFPR is used to define an output file for a printer.

M	DEVADDR=SYSxxx	Symbolic unit for the printer used for this file.
M	IOAREA1=xxxxxxxx	Name for the first output area.
O	ASOCFLE=xxxxxxxx	Name of the associated file for FUNC=RW, RPW, PW.
O	BLKSIZE=nnn	Length of one output area, in bytes. If omitted, 121 is assumed for 1403, 1443, 3203 or 3211; 136 is assumed for 3800 without TRC (or 137 with TRC): 64 is assumed for 2560 or 3525; 96 is assumed for 5203 or 5424/5425.[1]
O	CONTROL=YES	CNTRL macro used for this file. Omit CTLCHR for this file. Not allowed for 2560 or 5424/5425.
O	CTLCHR=xxx	(YES or ASA). Data records have control character. YES for S/370 character set; ASA for American National Standards Institute character set. Omit CONTROL for this file. Not allowed for 2560 or 5424/5425.
O	DEVICE=nnn	(1403, 1443, 2560P, 2560S, 3203, 3211, 3525, 3800, 5203. Specify 5425P or 5425S for 5424/5425 (P or S). Specify PRT1 for 3203-4, 3203-5, 3211, or 3289-4. If omitted, 1403 is assumed.[1]
O	ERROPT=xxxxxxxx	RETRY or the name of your error routine for 3211. IGNORE for 3525. Not allowed for other devices.[1]
O	FUNC=xxxx	(W, RW, RPW, PW) for 2560 or 5424/5425. (W[T], RW[T], RPW[T], PW[T] for 3525.
O	IOAREA2=xxxxxxxx	If two output areas are used, name of second area.
O	IOREG=(nn)	Register number, if two output areas used and PUT does not specify a work area. Omit WORKA.
O	MODNAME=xxxxxxxx	Name of PRMOD logic module for this DTF. If omitted, IOCS generates standard name. Not needed with 3800 advanced printer buffering.
O	PRINTOV=YES	PRTOV macro used for this file. Not allowed for 2560 or 5424/5425.
O	RDONLY=YES	Generate a read-only module. Requires a module save area for each task using the module.
O	RECFORM=xxxxxx	(FIXUNB, VARUNB, or UNDEF). If omitted, FIXUNB is assumed.
O	RECSIZE=(nn)	Register number if RECFORM=UNDEF.
O	SEPASMB=YES	DTFPR is to be assembled separately.
O	STLIST=YES	1403 selective tape listing feature is to be used.
O	TRC=YES	For 3800, output data lines include table reference character.
O	UCS=xxx	(ON) process data checks. (OFF) ignores data checks. Only for printers with the UCS feature, 3211, or 3800. If omitted, OFF is assumed.[1]
O	WORKA=YES	PUT specifies work area. Omit IOREG.

M=Mandatory
O=Optional

[1] 3211 remarks apply also to 3211-compatible printers (that is, with a device type code of PRT1).

Figure 3-40. DTFPR macro operands.

ASOCFLE=filename: This operand is used together with the FUNC operand to define associated files for the 2560, 3525, or 5424/5425. (For a description of associated files see *DOS/VSE Data Management Concepts*, as listed in the Preface.) ASOCFLE specifies the filename of an associated read and/or punch file, and enables macro sequence checking by the logic module of each associated file. One filename is required per DTF for associated files.

Figure 3-41 defines the filename specified by the ASOCFLE operand for each of the associated DTFs.

Code in FUNC= operand	filename specification in ASOCFLE= operand of		
	read DTFCD	punch DTFCD	print DTFPR
FUNC=RPW	filename of punch DTFCD	filename of print DTFPR	filename of read DTFCD
FUNC=PW		filename of print DTFPR	filename of punch DTFCD
FUNC=RW	filename of print DTFPR		filename of read DTFCD

Examples:
1. If FUNC=PW is specified,
   a.   specify the filename of the print DTFPR in the ASOCFLE operand of the punch DTFCD and
   b.   Specify the filename of the punch DTFCD in the ASOCFLE operand of the print DTFPR.
2. If FUNC=RPW is specified,
   a.   specify the filename of the punch DTFCD in the ASOCFLE operand of the read DTFCD, and
   b.   specify the filename of the print DTFPR in the ASOCFLE operand of the punch DTFCD, and
   c.   specify the filename of the read DTFCD in the ASOCFLE operand of the print DTFPR.

Figure 3-41. ASOCFLE operand usage with print associated files.

BLKSIZE=nnn: This operand specifies the length of IOAREA1. The maximum values which may be specified in this operand and the lengths assumed when it is omitted are given for the different devices in Figure 3-42.

CONTROL=YES: This operand is specified if the CNTRL macro will be issued for the file. If this operand is specified, omit CTLCHR. This operand is not allowed for the 2560 or 5424/5425.

CTLCHR= {YES|ASA}: This operand is specified if first-character control is used. The parameter ASA specifies the American National Standards Institute, Inc. character set. The entry CTLCHR=YES specifies the S/370 character set. If this parameter is specified, omit CONTROL. This operand must not be specified for the 2560 or 5424/5425.

If CTLCHR=ASA is specified for the 3525, the + character is not allowed. To print on the first line of a card, you must issue either a space 1 command or a skip to channel 1 command. For 3525 print associated files, you must issue a space 1 command to print on the first line of a card.

DEVADDR= {SYSLOG|SYSLST|SYSnnn}: This operand specifies the symbolic unit to be associated with the printer. SYSLOG and SYSLST must not be specified for the 2245, 2560, 3525, or 5424/5425.

Devices	Maximum length (in bytes) which can be specified [1]	Length assumed (in bytes) [2]
1403-1, -4	100	121
1403-6, -7	120	121
1403-2, -3, -5, -8, -9	132	121
1443	144	121
2560	384	64
3203	132	121
3203-4, -5	132	121
3211	150	121
3289-4	132	121
3525	64	64
3800	204 (without TRC)[3]	136 (without TRC)[3]
5203	132	96
5424/5425	128	96

[1]   RECFORM is FIXUNB or UNDEF and operand CTLCHR is not specified.

[2]   The parameter BLKSIZE=n is omitted.

[3]   For a 3800, the maximum length is 205 if TRC=YES is used, and the assumed length is 137.

Notes:
- If CTRCHR=YES/ASA is specified, add 1 byte to the maximum length which can be specified
- If RECFORM=VARUNB is specified add 4 bytes to the maximum value which can be specified.
- For the 2245, if RECFORM=VARUNB and CTLCHR=YES/ASA are specified, the maximum blocksize is 805 bytes.

Figure 3-42. Maximum and assumed lengths for the IOAREA1.

DEVICE= {1403|1443|2560P|2560S|3203|
|3211|3525|3800|5203|5425P|
5425S|PRT1}:
This operand specifies which device is used for the
file. The "P" and "S" included with the "2560" and
"5425" parameters specify primary or secondary
input hoppers. Specify 5425P/S for 5424(P/S).
"PRT1" refers to a 3211 or 3211-compatible printer.
If this operand is omitted, 1403 is assumed.

ERROPT={RETRY|IGNORE|name}: This ope-
rand specifies the action to be taken in the case of an
equipment check error. The functions of the param-
eters are described below.

RETRY can be specified only for a PRT1 printer.
RETRY indicates that if an equipment check with
command retry is encountered, the command is
retried once. If the retry is unsuccessful a message is
issued and the job canceled.

IGNORE can be specified only for the 3525.
IGNORE indicates that the error is to be ignored. The
address of the record in error is put in register 1 and
made available for processing. Byte 3, bit 3 of the
CCB is also set on (see Figure 4-2); you can check
this bit and take the appropriate action to recover
from the error. IGNORE must not be specified for
files with two I/O areas or a work area.

ERROPT=name can be specified only for a 3211-
compatible printer. It indicates that if an equipment
check with command retry is encountered, the com-
mand is retried once. If the retry is unsuccessful a
message is issued and the job canceled. With other
types of errors (for these see the CCB, Figure 4-2) an
error message is issued, error information is placed
in the CCB, and control is given to your error rou-
tine, where you may perform whatever actions are
desired. If any IOCS macros are issued in the rou-
tine, register 14 must be saved; if the operand
RDONLY=YES is specified, register 13 must also be
saved. To continue processing at the end of the rou-
tine, return to IOCS by branching to the address in
register 14.

FUNC={W[T]|RW[T]|RPW[T]|PW[T]}: This
operand specifies the type of file to be processed by
the 2560, 3525, or 5424/5425. W indicates print, R
indicates read, P indicates punch, and T (for the 3525
only) indicates an optional 2-line printer.

RW[T], RPW[T], and PW[T] are used,together with
the ASOCFLE operand, to specify associated files;
when one of these parameters, other than T, is speci-
fied for a printer file it must also be specified for the
associated file(s). Note: Do not use T for associated
files; it is valid only for printer files.

If a 2-line printer is not specified for the 3525,
multi-line print is assumed. T is ignored if CONTROL
or CTLCHR is specified.

IOAREA1=name: This operand specifies the name
of the output area.

IOAREA2=name: This operand specifies the name
of a second output area.

IOREG=(r): If two output areas and no work areas
are used, this operand specifies the register into
which IOCS will place the address of the area where
you can build a record. For (r) specify one of the
registers 2 to 12.

MODNAME=name: This operand may be used to
specify the name of the logic module that is used
with the DTF table to process the file. If the logic
module is assembled with the program, MODNAME
must specify the same name as the PRMOD macro. If
this operand is omitted, standard names are generat-
ed for calling the logic module. If two DTF macros
call for different functions that can be handled by a
single module, only one module is called.

PRINTOV=YES: This operand is specified if the
PRTOV macro is included in your program. This
operand is not allowed for the 2560 or 5424/5425.

RDONLY=YES: This operand is specified if the
DTF is used with a read-only module. Each time a
read-only module is entered, register 13 must con-
tain the address of a 72-byte doubleword-aligned
save area. Each task requires its own uniquely de-
fined save area. Each time an imperative macro
(except OPEN or OPENR) is issued, register 13 must
contain the address of the save area associated with
the task. The fact that the save areas are unique for
each task makes the module reentrant (that is, capa-
ble of being used concurrently by several tasks).

If an ERROPT routine issues I/O macros which use
the same read-only module that caused control to
pass to either error routine, your program must pro-
vide another save area. One save area is used for the
normal I/O, and the second for I/O operations in the
ERROPT routine. Before returning to the module
that entered the ERROPT routine, register 13 must be
set to the save area address originally specified for
the task.

If this operand is omitted, the module generated
is not reenterable and no save area need be estab-
lished.

RECFORM= {FIXUNB|UNDEF|VARUNB}:
The operand RECFORM=FIXUNB is specified when-
ever the record format is fixed. When the record
format is FIXUNB, this entry may be omitted.

The entry RECFORM=UNDEF is specified when-
ever the record format is undefined. If the output is
variable and unblocked, enter VARUNB.

RECSIZE=(r): This operand specifies the general
register (any one of 2 to 12) that will contain the
length of an output record of undefined format. The
length of each record must be loaded into the regis-
ter before issuing the PUT macro.

SEPASMB=YES: Include this operand only if the
DTFPR will be assembled separately. This causes a
CATALR card with the filename to be punched ahead
of the object deck and the filename to be defined as
an ENTRY point in the assembly. If the operand is
omitted, the assembler assumes that the DTF is being
assembled with the problem program and no
CATALR card is punched.

STLIST=YES: Include this operand if the selec-
tive tape listing feature (1403 only) is used. If this
entry is specified, the CONTROL, CTLCHR, and
PRINTOV entries are not valid and will be ignored if
specified. If this operand is specified, RECFORM
must have the parameter FIXUNB.

TRC=YES: This operand applies to the 3800 Print-
ing Subsystem; DEVICE=3800 should be specified.

TRC=YES specifies that a table reference character is
included as the first byte of each output data line
(following the optional print control character). The
printer uses the table reference character to select
the character arrangement table corresponding to
the order in which the table names have been speci-
fied with the CHAR parameter on the SETPRT job
control statement (or SETPRT macro instruction).

If a printer other than a 3800 is specified on the
DEVICE parameter, any table reference character
sent to that printer is treated as data.

UCS= {OFF|ON}: For a printer with the universal
character set feature, or for a 3800 Printing Subsys-
tem, this operand determines whether data checks
occurring in case of unprintable characters are indi-
cated to the operator or printed as blanks. The entry
is especially useful if you are using first-character
forms control and have modules that cannot process
the CNTRL macro.

ON   Data checks are processed with an operator
     indication.

OFF Data checks are ignored and blanks are print-
     ed for the unprintable character.

WORKA=YES: If output records are processed in
work areas instead of in the I/O areas, specify this
operand. You must set up the work area in storage.
The address of the work area, or a general-purpose
register which contains the address, must be speci-
fied in each PUT macro.

PRMOD Macro:

Listed here are the operands you can specify for PRMOD. The first card contains PRMOD in the operation field and may contain a module name in the name field.

If advanced printer buffering is used on your 3800 Printer Subsystem, the PRMOD macro is not needed.

CONTROL=YES: Include this operand if CNTRL macros are used with the associated DTFs. The module also processes files that do not use the CNTRL macro. If CONTROL is specified, the CTLCHR operand must not be specified.

The CONTROL operand is not allowed for the 2560 or 5424/5425.

CTLCHR={YES|ASA}: Include this operand if first-character carriage control is used. Any DTF used with the module must have the same operand. If CTLCHR is specified, CONTROL must not be specified.

CTLCHR must not be specified for the 2560 or 5424/5425.

If CTLCHR=ASA is specified for the 3525, the + character is not allowed. For 3525 print (not associated) files, you must issue either a space 1 command or skip to channel 1 command to print on the first line of a card. For 3525 print associated files, you must issue a space 1 command to print on the first line of a card.

If CTLCHR=ASA and RDONLY=YES are specified in a multitasking environment where more than one DTFPR uses the same module, overprinting may occur.

DEVICE={1403|1443|2560P|2560S|
       3203|3211|3525|3800|5203|5425P|
       5425S|PRT1}:
This operand specifies which device is used for the file. The "P" and "S" included with the "2560" and "5425" parameters specify primary or secondary input hoppers; regardless of which is specified, however, the module generated will handle DTFs specifying either hopper. Specify 5425P/S for 5424P/S.

Any DTF to be used with this module must have the same operand (except as just noted concerning the "P" and "S" specification for the 2560 or 5424/5425).

ERROPT=YES: This operand must be specified if ERROPT=name is specified in a DTFPR that is to be used with the module. (ERROPT=name is applicable to a 3211-compatible printer only.) If ERROPT is not specified in the DTFPR, or if ERROPT=RETRY (3211) or ERROPT=IGNORE (3525) is specified, ERROPT=YES must be omitted.

FUNC={W[T]|RW[T]|RPW[T]|PW[T]}: This operand specifies the type of file to be processed by the 2560, 3525, or 5424/5425. Any DTF used with the module must include the same operand. W indicates print, R indicates read, P indicates punch, and T (for the 3525 only) indicates an optional 2-line printer.

RW[T], RPW[T], and PW[T] are used to specify associated files; when one of these parameters is specified for a printer file it must also be specified for the associated file(s).

If a 2-line printer is not specified for the 3525, multi-line print is assumed. T is ignored if CONTROL or CTLCHR is specified.

IOAREA2=YES Include this operand if a second I/O area is used. Any DTF used with the module must also include the IOAREA2 operand.

PRINTOV=YES Include this operand if PRTOV macros are used with the associated DTFs. The module also processes any files that do not use the PRTOV macro.

This operand is not allowed for the 2560 or 5424/5425.

RDONLY=YES: This operand causes a read-only module to be generated. Whenever this operand is specified, any DTF used with the module must have the same operand.

RECFORM={FIXUNB|VARUNB|UNDEF}: This operand causes a module to be generated that processes the specified record format: fixed-length, variable-length, or undefined. Any DTF used with the module must include the same operand.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and defines the module name as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

STLIST=YES: Include this operand if the selective tape listing feature (1403 only) is used. If this entry is specified, the CONTROL, CTLCHR, and PRINTOV entries are not valid, and are ignored if

supplied. If this operand is specified, RECFORM must specify FIXUNB.

TRC=YES: Include this operand to specify whether the module is to test the TRC bit in the DTFPR or iqnore that bit. If TRC=YES is specified, the generated module can process output files with table reference characters and those without.

WORKA=YES: Include this operand if records are processed in work areas instead of in I/O areas. Any DTF used with the module must have the same operand.

Standard PRMOD Names

Each name begins with a 3-character prefix (IJD) followed by a 5-character field corresponding to the options permitted in the generation of the module.

PRMOD name = IJDabcde

a = F RECFORM=FIXUNB
  = V RECFORM=VARUNB
  = U RECFORM=UNDEF

b = A CTLCHR=ASA
  = Y CTLCHR=YES
  = C CONTROL=YES
  = S STLIST=YES
  = Z none of these is specified
  = T DEVICE=3525 with 2-line printer
  = U DEVICE=2560
  = V DEVICE=5425

c = B ERROPT=YES and PRINTOV=YES
  = P PRINTOV=YES, DEVICE is not 3525, and ERROPT is not specified
  = 1 PRINTOV=YES, DEVICE=3525, and FUNC=W[T] or omitted
  = F PRINTOV=YES, DEVICE=3525, and FUNC=RW[T]
  = C PRINTOV=YES, DEVICE=3525, and FUNC=PW[T]
  = D PRINTOV=YES, DEVICE=3525, and FUNC=RPW[T]
  = Z neither PRINTOV nor ERROPT is specified, and DEVICE is not a 3525
  = O PRINTOV=YES not specified, DEVICE=3525, and FUNC=W[T] or omitted
  = R PRINTOV=YES not specified, DEVICE=3525, and FUNC=RW[T]
  = S PRINTOV=YES not specified, DEVICE=3525, and FUNC=PW[T]
  = T PRINTOV=YES not specified, DEVICE=3525, and FUNC=RPW[T]
  = E ERROPT=YES and PRINTOV=YES is not specified
  = U FUNC=W or omitted and DEVICE=2560 or 5425
  = V FUNC=RW and DEVICE=2560 or 5425

    = W FUNC=PW and DEVICE=2560 or 5425
    = X FUNC=RPW and DEVICE=2560 or 5425

d = I IOAREA2=YES
  = Z IOAREA2=YES is not specified

e = V RDONLY=YES and WORKA=YES
  = W WORKA=YES
  = Y RDONLY=YES
  = Z neither is specified

Subset/Superset PRMOD Names

Figure 3-43 shows the subsetting and supersetting allowed for PRMOD names. Two of the operands allow subsetting. For example, the module name IJDFCPIW is a superset of the module names IJDFCZIW and IJDFZZIW. No subsetting or supersetting of PRMOD names is allowed for the 2560 or 5424/5425.

The IBM-supplies preassembled logic modules do not have TRC=YES. The system programmer can reassemble them with TRC=YES to support 3800 table reference characters. Although the code that is generated for a module assembled with TRC=YES is different from the code that is generated for a module with TRC=NO, the module name is the same. If some, but not all PRMOD logic modules are reassembled this way, it may interfere with subsetting or supersetting.

```
        *  *  *  *  *
I  J  D  F  A  U  I  V
         V  Y  V  Z  W
         U  S  W     Y
         T  X     Z
         U  +
         V  P
         +  Z
         C  +
         Z  I
            O
            +
            F
            R
            +
            +
            C
            S
            +
            D
            T
            +
            B
            E
```

+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.
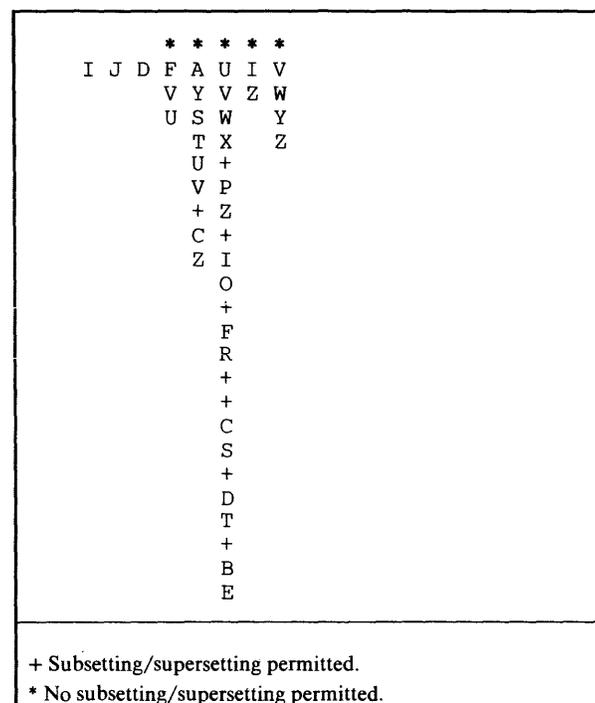
Figure 3-43. Subsetting and supersetting of PRMOD names.

DTFPT Macro

A DTF entry is included for every paper tape input or output file that is processed by the program. The characteristics of a paper tape file are given in the *DOS/VSE Macro User's Guide*, as listed in the Preface.

Applies to				
Input	Output			
x	x	M	BLKSIZE=n	Length of your I/O areas.
x	x	M	DEVADDR=SYSnnn	Symbolic unit to be associated with this file.
x	x	M	IOAREA1=xxxxxxxx	Name of first I/O area.
x		O	EOFADDR=xxxxxxxx	Name of your end-of-file routine.
	x	O	DELCHAR=X'nn'	Delete character.
x	x	O	DEVICE=nnnn	(2671, 1017, 1018). If omitted, 2671 is assumed.
	x	O	EORCHAR=X'nn'	End-of-record character. (For RECFORM=UNDEF).
x	x	O	ERROPT=xxxxxxxx	(IGNORE, SKIP, or error routine name). Prevents job termination on error records.
	x	O	FSCAN=xxxxxxxx	(For shifted codes). Name of your scan table used to select figure groups.
x		O	FTRANS=xxxxxxxx	(For shifted codes). Symbolic address of your figure shift translate table.
x	x	O	IOAREA2=xxxxxxxx	Name of second I/O area.
x	x	O	IOREG=(nn)	Used with two I/O areas. Register (2-12) containing current record address.
	x	O	LSCAN=xxxxxxxx	(For shifted codes). Name of your scan table used to select letter groups.
x		O	LTRANS=xxxxxxxx	(For shifted codes). Name of your letter shift translate table.
x	x	O	MODNAME=xxxxxxxx	For module names other than standard.
x	x	O	OVBLKSZ=n	Used if I/O records are compressed or expanded.
x	x	O	RECFORM=xxxxxx	(FIXUNB or UNDEF). If omitted, FIXUNB is assumed.
x	x	O	RECSIZE=(nn)	Register containing the record length.
x		O	SCAN=xxxxxxxx	Name of your scan table for shift or delete character.
x	x	O	SEPASMB=YES	DTF is assembled separately.
x	x	O	TRANS=xxxxxxxx	Name of your table for code translation.
x		O	WLRERR=xxxxxxxx	Name of wrong-length-record error routine.

M=Mandatory
O=Optional

Figure 3-44. DTFPT macro operands.

BLKSIZE=n: This operand specifies the length of the input or output area. The maximum block size is 32,767 bytes.

DELCHAR=X'nn': This operand specifies the configuration of the delete character and must be used for output files only, that is, when DEVICE=1018 is specified. The constant X'nn' consists of two hexadecimal digits. The delete character is used in the error recovery procedure, and you must specify the correct configuration in accordance with the number of tracks of the output tape, as follows:

    X'1F' for five tracks.
    X'3F' for six tracks.
    X'7F' for seven tracks.
    X'FF' for eight tracks.

Note: The delete character is required only if the 1018 has the error correction feature.

DEVADDR=SYSnnn: This operand specifies the logical unit (SYSnnn) associated with this file. An actual channel and unit are assigned to the unit by an ASSGN job control statement. The ASSGN statement contains the same symbolic name as DEVADDR.

DEVICE= {2671|1017|1018}: This operand is required only to specify the paper tape I/O device. If this entry is omitted, 2671 is assumed.

EOFADDR=name: This operand specifies the name of your end-of-file routine. IOCS automatical-

ly branches to this routine on an end-of-file condition if the end-of-file switch is set on. The routine can execute any operation required for the end-of-file, issue the CLOSE or CLOSER macro for the file, or return to IOCS by branching to the address in register 14. In the latter case, IOCS reads in the next record. The end-of-file condition cannot occur on the 1018.

EORCHAR=X'nn': This operand specifies the user-defined end-of-record (EOR) character, where nn is two hexadecimal digits. It must be used for output files with undefined record format only. IOCS writes this character after the last character of the undefined record.

ERROPT= {IGNORE|SKIP|name}: This operand is specified if you do not want a job terminated when the standard recovery procedure cannot recover from a read or write error. If the ERROPT entry is omitted and a read or write error occurs, IOCS terminates the job.

For input files, IGNORE allows IOCS to handle the record as if no errors were detected. If SKIP is specified, IOCS skips the record in error and reads the next record.

For output files with shifted codes, ERROPT cannot be specified. For unshifted codes, the options ERROPT=IGNORE and ERROPT=name can be specified. IGNORE allows IOCS to handle the record as if no errors were detected.

The ERROPT=SKIP option is ignored and causes IOCS to terminate the job.

If two I/O areas are used, the CLOSE or CLOSER macro checks the last record, and the option ERROPT=name is treated as option ERROPT=IGNORE.

For name, specify the symbolic address of your error routine that will process errors. On an error condition, IOCS reads or writes the complete record, including the error character(s), and then branches to the error routine. At the end of the error routine, return to IOCS by branching to the address in register 14. The next record is then read or written. You must not issue any GET or PUT macros for records in the error block. If the error routine contains any other IOCS macros, the contents of register 14 must be saved and restored.

FSCAN=name: This operand must be included for every output file using a shifted code. Omit this operand for an input file. The operand specifies the name of a scan table in your program used to select groups of figures. This table must conform to the

specifications of the machine instruction TRT. The entry in the table for each letter character must be the letter shift character, and all other entries must be hexadecimal zero. Any deviation from this results in incorrect translation.

FTRANS=name: This operand must be included for every input file using a shifted code and is not permitted for output files. It specifies the name of a figure shift table in your program. This table must conform to the specifications of the machine instruction TR.

IOAREA1=name: This operand specifies the name of an input or output area.

IOAREA2=name: This operand specifies the name of a second input or output area. When this operand is specified, IOCS overlaps the I/O operation in one area with the processing of the record in the other.

IOREG=(r): This operand must be included if two input or output areas are used. For input, it specifies the register into which IOCS puts the address of the logical record available for processing. For output, it specifies the register that contains the address of the area in which your program can build a record. Any register from 2 to 12 may be specified.

LSCAN=name: This operand must be included for every output file using a shifted code and is not permitted for input files. It specifies the name of a scan table in your program used to select groups of letters. This table must conform to the specifications of the machine instruction TRT. The entry in the table for each figure character must be the figure shift character, and all other entries must be hexadecimal zero. Any deviation from this results in incorrect translation.

LTRANS=name: This operand must be included for every input file using a shifted code and is not permitted for output files. It specifies the name of a letter shift table in your program. This table must conform to the specifications of the machine instruction TR.

MODNAME=name: This operand specifies the name of the logic module used with the DTF table to process the file. If the logic module is assembled with the program, the MODNAME operand in this DTF must specify the same name as the PTMOD macro. If the operand is omitted, IOCS generates standard names for calling the logic module.

OVBLKSZ=n: For input files, this operand specifies the number of characters to be read (before translation and compression) to produce the number of characters specified in the BLKSIZE entry. OVBLKSZ is used only when SCAN=name and RECFORM=FIXUNB are both specified. If OVBLKSZ is omitted, IOCS assumes that the number of characters to be read is equal to the number specified in the BLKSIZE entry. The maximum value is 32,767 bytes.

For output files, OVBLKSZ specifies the number of characters indicated in the BLKSIZE entry, plus the number of shift characters to be inserted. If the size of OVBLKSZ is large enough to allow the insertion of all the shift characters required to build the output record, a single WRITE operation results from a PUT macro. On the other hand, if the size of OVBLKSZ (which must be at least one position larger than BLKSIZE) does not permit the insertion of all the shift characters, several WRITE operations result from a PUT macro. OVBLKSZ is used only when LSCAN and FSCAN are specified with the FIXUNB format. If OVBLKSZ is specified with UNDEF format, it is ignored.

RECFORM= {FIXUNB|UNDEF}**:** This operand specifies the record format for the file. Specify either format for shifted or unshifted codes. If the record format is FIXUNB, this entry may be omitted.

RECSIZE=(r): This operand specifies the number of a register (any one of 2 to 12) that contains the length of the input or output record. This entry is optional for input files. If present, IOCS loads the length of each record read into the specified register. If input files contain shift codes or other characters requiring deletion, IOCS loads the compressed record length into the specified register.

For output files, this entry must be included for undefined records. Before translation, your program must load each record length into the designated register before issuing the PUT macro for the record.

SCAN=name: This operand must be included for all input files using shifted codes. It may also be included if you wish to delete certain characters from each record. The SCAN entry specifies the symbolic name of a table provided by your program. This table must conform to the specifications of the machine instruction TRT. It must contain nonzero entries for all delete characters and, where appropriate, for the figure and letter shift characters.

The table entry for the figure shift character must be X'04'; for the letter shift character, the entry must be X'08'; delete entries must be X'0C'. All other entries

in the table must be X'00'. Otherwise, incorrect translation results and a program check may occur.

The table must be large enough to hold the maximum value of coding for the tape being processed; that is, 255 bytes for 8-track tape. This prohibits erroneous coding on the tape from causing a scan function beyond the limits of the scan table.

SEPASMB=YES: Include this operand only if the DTFPT is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TRANS=name: The TRANS operand specifies the symbolic name of a table provided within your program. This table must conform to the specifications of the machine instruction TR. For input files, include this entry if a nonshifted code is to be translated into internal system code. Omit the FTRANS and LTRANS entries if this entry is present. If none of these three entries is present, no translation takes place. For output files, include this entry if the internal system code is translated into a shifted or nonshifted code, depending on whether the FSCAN and LSCAN entries are present or omitted.

WLRERR=name: This operand applies only to paper tape input files when RECFORM=UNDEF is specified.

When IOCS finds a wrong-length record, it branches to the symbolic name specified in the WLRERR entry. If this entry is omitted and the ERROPT entry is included, IOCS considers the error uncorrectable and uses the ERROPT option specified. Absence of both ERROPT and WLRERR entries causes the wrong-length record to be accepted as a normal record. Wrong-length checking is not performed for fixed-length records because a fixed number of characters is read in each time. IOCS detects overlength undefined records when the incoming record fills the input area. The input area must, therefore, be at least one position longer than the longest record anticipated.

At the end of the WLRERR routine, return to IOCS by branching to the address in register 14. The next IOCS read operation will normally cause the remainder of the overlength, undefined record to be read. If any other IOCS macros are included in the record-length error routine, the contents of register 14 must be saved and restored.

Note: A wrong-length condition appears during the first read operation on a 1017 if the combined length of the tape leader and the first record is greater than the length of the longest record anticipated (the length specified in BLKSIZE).

PTMOD Macro

Listed here are the operands you can specify for PTMOD; Figure 3-45 shows the only possible combination of these operands and describes the resultant modules.

Operand *				Resulting Module
DEVICE=	RECFORM=	SCAN=	TRANS=	
2671 **	FIXUNB **			Does not handle translation or shift or delete characters
2671 **	FIXUNB **		YES	Handles translation of unshifted codes, but not delete characters
2671 **	FIXUNB **	YES		Handles shift and delete characters for records of fixed un-blocked format
2671 **	FIXUNB **	YES		
2671 **	UNDEF	YES		Handles shift and delete characters for records of undefined format
1017	FIXUNB **			Does not handle translation or shift or delete characters
1017	FIXUNB **		YES	Handles translation of unshifted codes, but no delete characters
1017	FIXUNB **	YES		Handles shift and delete characters for records of fixed un-blocked format
1017	UNDEF	YES		Handles shift and delete characters for records of undefined format
1018	FIXUNB **		YES	Handles translation of unshifted codes, if specified in DTFPT, for records of fixed unblocked format
1018	FIXUNB **			
1018	UNDEF			Handles translation of unshifted codes, if specified in DTFPT, for records of undefined format
1018	UNDEF		YES	
1018	FIXUNB **	YES		Handles shift characters for records of fixed unblocked format
1018	UNDEF	YES		Handles shift characters for records of undefined format
* In all cases, SEPASMB=YES may either be specified or omitted				
** Specified explicitly or by default				

Figure 3-45. PTMOD operand combinations.

DEVICE= {2671|1017|1018}: Required only to specify an I/O device other than 2671 used by the module. Any DTF used with the module must have the same operand. 2671 is assumed if this operand is omitted.

RECFORM= {FIXUNB|UNDEF}: Required only if the operand SCAN=YES is present. If records of undefined format using the SCAN option are translated, specify the UNDEF parameter. If records of fixed unblocked format are translated, the FIXUNB parameter may be specified or omitted.

SCAN=YES: Required for records containing shift characters and/or characters that are automatically deleted by IOCS.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and defines the module name as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assem-

bled with the DTF and the problem program and no CATALR card is punched.

TRANS=YES: Required only if records using an unshifted code are translated and if the operand SCAN=YES is not specified.

Standard PTMOD Names

Each name begins with a 3-character prefix (IJE) and continues with a 5-character field corresponding to the options permitted in the generation of the module.

PTMOD name =IJEabcde

a = S SCAN=YES
  = Z SCAN=YES is not specified

b = T TRANS=YES (SCAN=YES is not specified)
  = Z TRANS=YES is not specified

c = F RECFORM=FIXUNB, and SCAN=YES
  = U RECFORM=UNDEF, and SCAN=YES
  = Z SCAN=YES is not specified, and/or DEVICE=1018

d = 1 DEVICE=1017
  = 2 DEVICE=1018

= Z DEVICE=2671, or if this entry is omitted

e = Z always

Subset/Superset PTMOD Names

Figure 3-46 shows the PTMOD names. No subsetting or supersetting is allowed.

```
          *  *  *  *
I  J  E   Z  Z  Z  Z  Z
          Z  T  Z  Z
          S  Z  F  Z
          S  Z  U  Z
          Z  Z  Z  1
          Z  T  Z  1
          S  Z  F  1
          S  Z  U  1
          S  Z  Z  2
          Z  T  Z  2

* No subsetting/supersetting permitted.
```

Figure 3-46. Subsetting and supersetting of PTMOD names.

DTFSD Macro

The DTFSD macro defines sequential (consecutive) processing for a file contained on a DASD. Only IBM standard label formats are processed.

Applies to					
Input	Output	Work			
x	x	x	M	BLKSIZE=nnnn	Length of one I/O area, in bytes
x		x	M	EOFADDR=xxxxxxxx	Name of your end-of-file routine
x	x		M/O	IOAREA1=xxxxxxx	Name of first I/O area. Optional for some files in Control Interval format.
x	x	x	O	CISIZE=nnnnn	Size of FBA Control Interval. If omitted and DEVICE=FBA, default is 0.
x	x	x	O	CONTROL=YES	CNTRL macro is used for this file
		x	O	DELETFL=NO	CLOSE, CLOSER macro is not to delete format-1 and format-3 labels for work file
x	x	x	O	DEVADDR=SYSnnn	Symbolic unit required only when not provided on an EXTENT statement
x	x	x	O	DEVICE=nnnn	(2311, 2314, 3330, 3340, 3350, FBA). Specify any device for a 3330-11 or 3350. Specify FBA for 3310 or 3370. If omitted, 2311 is assumed.
x	x	x	O	ERREXT=YES	Additional error and ERET are desired. Specify ERROPT also
x	x	x	O	ERROPT=xxxxxxxx	(IGNORE, SKIP, or name of error routine). Prevents job termination on error records. Do not use SKIP for output files
x	x		O	FEOVD=YES	Forced end of volume for disk is desired
x		x	O	HOLD=YES	Employ the track hold function
x	x		O	IOAREA2=xxxxxxx	If two I/O areas are used, name of second area
x	x		O	IOREG=(nn)	Register number. Use only if GET or PUT does not specify work area or if two I/O areas are used. Omit WORKA
x	x		O	LABADDR=xxxxxxx	Name of your routine to check/write user-standard labels
x	x	x	O	MODNAME=xxxxxxxx	Name of SDMODxx logic module for this DTF. If omitted, IOCS generates standard name. Ignored if DEVICE=FBA is specified.
		x	O	NOTEPNT=xxxxxxx	(YES or POINTRW). YES if NOTE/POINTR/POINTW/POINTS used. POINTRW if only NOTE/POINTR/POINTW used
	x	x	O	PWRITE=YES	For FBA files only, specify for a physical write of each logical block.
x	x	x	O	RDONLY=YES	Generates a read-only module. Requires a module save area for each task using the module
x	x	x	O	RECFORM=xxxxxx	(FIXUNB, FIXBLK, VARUNB, VARBLK, SPNUNB, SPNBLK, or UNDEF). For work files use FIXUNB or UNDEF. If omitted, FIXUNB is assumed

M=Mandatory
O=Optional

Figure 3-47. DTFSD macro operands (Part 1 of 2).

Applies to					
Input	Output	Work			
x	x		O	RECSIZE=nnnnn	If RECFORM=FIXBLK, number of characters in record. If RECFORM=SPNUNB, SPNBLK, or UNDEF, register number. Not required for other records
x	x		O	SEPASMB=YES	DTFSD is to be assembled separately.
x	x		O	TRUNCS=YES	RECFORM=FIXBLK or TRUNC macro used for this file
x	x	x	O	TYPEFLE=xxxxxx	(INPUT, OUTPUT, or WORK). If omitted, INPUT is assumed
x		x	O	UPDATE=YES	Input file or work file is to be updated
	x		O	VARBLD=(nn)	Register number if RECFORM=VARBLK and records are built in the output area. Omit if WORKA=YES
	x	x	O	VERIFY=YES	Check disk records after they are written.
x			O	WLRERR=xxxxxxxx	Name of your wrong-length-record routine
x	x		O	WORKA=YES	GET or PUT specifies work area. Omit IOREG. Required for RECFORM=SPNUNB or SPNBLK.

M=Mandatory
O=Optional

Figure 3-47. DTFSD macro operands (Part 2 of 2).

BLKSIZE=n: Enter the length of the I/O area. If the record format is variable or undefined, enter the length of the I/O area needed for the largest block of records.

For output files, the first 8 bytes of the I/O area (whose address is specified in the IOAREA1 operand) must be allotted for IOCS to construct a count field.

The BLKSIZE parameter on the DLBL statement overrides the DTFSD BLKSIZE specification if the device assigned is a 3330-11 or 3350, and if RECFORM=xxxBLK. For an *output* file, the records are blocked according to the size specified by the appropriate BLKSIZE parameter (from the DLBL statement if it was specified; otherwise from the DTFSD). For an *input* file, the BLKSIZE specification must match the format of the data as it resides on the disk.

To use the DLBL BLKSIZE parameter, your program must:

- Run on a system with RPS support sysgened. The logic module must be in the SVA, or there must be enough space available in the SVA to load the logic module.

- Have GETVIS space for an RPS DTF extension and new buffers.

- Specify DTFSD RECFORM=xxxBLK.

If there is no RPS support, if the file is not on a 3330-11 or 3350, if the GETVIS fails, or if the file is not a blocked data file, the job is canceled. For blocked files with fixed length records, BLKSIZE must be a multiple of RECSIZE or the job will be canceled.

When DEVICE=FBA, or if CISIZE=n is specified, BLKSIZE determines logical block size. For FBA DASD, maximum value is 32761 (that is, seven bytes less than maximum CISIZE). The BLKSIZE value for output files must include eight bytes for a count area to provide compatibility between FBA and CKD DASD.

CISIZE=n: This operand specifies the control interval size for an FBA device assigned to a non-system file logical unit. If assigned to a system file (SYSRDR, SYSIPT, SYSLST, or SYSPCH), the operand is ignored. The value n must be a multiple of the FBA block size and, if greater than 8K, must be a multiple of 2K. The maximum value is 32768 (32K) except when assigned to SYSLST or SYSPCH, when the maximum is 30720 (30K).

If DEVICE=FBA is specified, and CISIZE is omitted, CISIZE=0 is assumed. Control Interval size may be overridden for an output file at execution time by specifying the CISIZE parameter on the DLBL control statement. For an input file, the CISIZE value in the format-1 label is used.

CONTROL=YES: This operand is specified if a CNTRL macro is to be issued for the file. A CCW is generated for control commands. For an FBA file, this operand is ignored.

DELETFL=NO: Specify this operand if the CLOSE or CLOSER macro is not to delete the format-1 and format-3 label for a work file. The operand applies to work files only.

DEVADDR=SYSnnn: This operand must specify the symbolic unit associated with the file if an extent is not provided. A job control EXTENT statement is not required for single-volume input files. If an EXTENT statement is provided, its specification overrides any DEVADDR specification. SYSnnn represents an actual I/O address, and is used in the ASSGN job control statement to assign the actual I/O device address to this file.

DEVICE= {2311|2314|3330|
                3340|3350|FBA}:
This operand is included to specify the device on which the file is located. If no device is specified, 2311 is assumed. If an FBA device is ASSGNed to the file, DEVICE= is ignored.

For devices supported by DOS/VSE and not included in the above operand specification, specify device codes as listed in Figure 3-48.

EOFADDR=name: This operand specifies the name of your end-of-file routine. IOCS automatically branches to this routine on an end-of-file condition. In this routine, you can perform any operations required at end of file (you generally issue the CLOSE or CLOSER macro).

ERREXT=YES: This operand enables your ERROPT or WLRERR routine to return to SDMOD by means of the ERET macro. It also enables unrecoverable I/O errors (such as "record not found") occuring before a data transfer takes place to be indicated to your program. For ERREXT facilities, the ERROPT operand must be specified. To take full advantage of this option, code the ERROPT=name parameter.

ERREXT=YES is assumed if DEVICE=FBA is specified, or if an FBA device is ASSGNed to the file.

ERROPT= {IGNORE|SKIP|name}: This operand is specified if a job is not to be terminated when a read or write error cannot be corrected in the disk error routines. The disk error routines normally retry failing I/O operations several times before considering the error unrecoverable. Once the error is considered unrecoverable, the job is terminated unless the ERROPT operand is specified. The functions of the parameters are explained below.

IGNORE
> The error condition is ignored. The records are made available for processing. When reading spanned records, the whole spanned record or block of spanned records is returned, rather than just the one physical record in which the error occurred.

> On output, the physical record or control interval in which the error occurred is ignored as if it were written correctly. If possible, any remaining spanned record segments are written.

SKIP
> No records in the error block or control interval are made available for processing. The next block or control interval is read from the disk, and processing continues with the first record of that block. When reading spanned records, the whole spanned record or block of spanned records is skipped, rather than just one physical record.

> On an UPDATE=YES file, the physical record or control interval in which the error occurred is ignored as if it were written correctly. If possible, any remaining spanned record segments are written.

DEVICE = specification	Device to be ASSGNed								
	2311	2314	2319	3330-1,2*	3330-11**	3340, 35MB	3340, 70MB	3350	3310, 3370
Default	x				x			x	x
2311	x				x			x	x
2314		x	x		x			x	x
3330				x	x			x	x
3340					x	x	x	x	x
3350					x			x	x
FBA									x

* Also 3350 in 3330-1 compatibility mode.
** Also 3350 in 3330-11 compatibility mode.

Figure 3-48. DEVICE= specifications for DTFSD.

name
> IOCS branches to your error routine named by this parameter. In this routine you can process or make note of the error condition as desired.

FEOVD=YES: This operand is specified if a forced end of volume for disk feature is desired. It forces the end-of-volume condition before physical end of volume occurs. When the FEOVD macro is issued, the current volume is closed, and I/O processing continues on the next volume.

HOLD=YES: This operand may be specified only if the track hold function was specified at system generation time and if it is employed when a data file or a work file is referenced for updating.

IOAREA1=name: This operand specifies the symbolic name of the I/O area used by the file. IOCS either reads or writes records using this area. If DEVICE=FBA is specified, this operand is not mandatory. It is ignored if WORKA=YES is specified. It is optional for input files if IOREG=(r) is specified. It is also optional for output files with fixed length records without truncation if IOREG=(r) is specified, but is required for all other FBA files.

For variable-length or undefined records, this area must be large enough to contain the largest block or record.

IOAREA2=name: If two I/O areas are used by GET or PUT, this operand is specified. When variable length records are processed, the size of the I/O area must include four bytes for the block size. For output files, the I/O area must include eight bytes to build a count field.

IOREG=(r): This operand specifies the general purpose register (any of 2 to 12) in which IOCS puts the address of the logical record that is available for processing. At OPEN time, for output files, IOCS puts into the register specified the address of the area where you can build a record. The same register may be used for two or more files in the same program, if desired. If this is done, the program must store the address supplied by IOCS for each record.

This operand must be specified if blocked input or output records are processed in one I/O area, or if two I/O areas are used and the records are processed in both I/O areas.

For an FBA file, if IOAREA(s) are not specified, the register specified by IOREG will point directly to data in the control interval buffer.

LABADDR=name: Specifies the name of the rou-

tine in which you process your own labels.

MODNAME=name: This operand may be used to specify the name of the logic module that will be used with the DTF table to process the file. If the logic module is assembled with the program, MODNAME must specify the same name as the SDMODxx macro.

This operand is ignored if DEVICE=FBA is specified. If this operand is omitted for other than FBA files, standard names are generated for calling the logic module. If two DTF macros call for different functions that can be handled by a single module, only one module is called.

NOTEPNT= {POINTRW|YES}: The parameter POINTRW is specified if a NOTE, POINTR, or POINTW macro is issued for the file. If the parameter YES is specified, NOTE, POINTR, POINTW, and POINTS macros may be issued for the file.

PWRITE=YES: This operand is specified if formatting output operations (PUT for data files or WRITE SQ for work files) are to cause a physical write for each logical block. If omitted, the physical write takes place only when the control interval buffer is full.

RDONLY=YES: This operand is specified if the DTF is used with a read-only module. Each time a read-only module is entered, register 13 must contain the address of a 72-byte doubleword-aligned save area. Each task should have its own uniquely defined save area. When an imperative macro (except OPEN, OPENR, or LBRET) is issued, register 13 must contain the address of the save area associated with the task. The fact that the save areas are unique for each task makes the module reentrant (that is, capable of being used concurrently by several tasks).

If an ERROPT or WLRERR routine issues I/O macros using the same read-only module that caused control to pass to either error routine, your program must provide another save area. One save area is used for the normal I/O operations, and the second for I/O in the ERROPT or WLRERR routine. Before returning to the module that entered the error routine, register 13 must be set to the save area address originally specified for the task.

As all control interval format logic modules are reentrant and read-only, this operand is ignored if DEVICE=FBA is specified. As the control interval format and RPS logic modules support DTFs with either read-only or non-read-only modules, no save area need be established.

If the operand is omitted for other than FBA files, the module generated is not reenterable and no save area need be established.

RECFORM= {FIXUNB|FIXBLK|VARUNB|
**   VARBLK|SPNUNB|SPNBLK|**
**   UNDEF}:**
This operand specifies the type of records for input or output. Enter one of the following parameters:

FIXUNB   For fixed-length unblocked records
FIXBLK   For fixed-length blocked records
VARUNB   For variable-length unblocked records
VARBLK   For variable-length blocked records
SPNUNB   For spanned variable-length unblocked records
SPNBLK   For spanned variable-length blocked records
UNDEF   For undefined records

If RECFORM=SPNUNB or RECFORM=SPNBLK is specified and RECSIZE=(r) is not specified, an assembler diagnostic (MNOTE) is issued, and register 2 is assumed. If WORKA=YES is omitted, an MNOTE is issued and WORKA=YES is assumed. If RECFORM is omitted, FIXUNB is assumed.

RECSIZE= {n|(r)}: For fixed-length blocked records, RECSIZE is required. It specifies the number of characters in each record.

Register notation must be used when processing spanned or undefined records. When processing undefined records and variable-length spanned records, RECSIZE is required for output files and is optional for input files. The operand is invalid for work files. It specifies a general register (any one of 2 to 12) that contains the length of the record. On output, you must load the length of each record into the designated register before issuing a PUT macro. If specified for input, IOCS provides the length of the record transferred to virtual storage.

SEPASMB=YES: Include this operand only if the DTFSD is assembled separately. This causes a CATALR card with the filename to be punched ahead of the object deck and the filename to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the DTF is being assembled with the problem program and no CATALR card is punched.

TRUNCS=YES: This operand is specified if FIXBLK DASD files contain short blocks embedded within an input file or if the input file was created with a module that specified TRUNCS. This entry is also specified if the TRUNC macro is issued for a FIXBLK output file.

TYPEFLE= {INPUT|OUTPUT|WORK}: Use this operand to indicate whether the file is an input or output file. If WORK is specified, a work file is to be used. If INPUT or OUTPUT is specified, the GET or PUT macros, respectively, can be used. If WORK is specified, the READ and WRITE, NOTE and POINTx, and CHECK macros can be used.

UPDATE=YES: This operand must be included if the DASD input or work file is updated – that is, if disk records are read, processed, and then re-written in the same disk record locations from which they were read. CLOSE writes any remaining records in sequence onto the disk.

This operand is invalid for a file on an FBA DASD assigned to a system logical unit (SYSRDR, SYSIPT, SYSLST, or SYSPCH). If a PUT is attempted to an input file, the job will be terminated.

VARBLD=(r): Whenever variable-length blocked records are built directly in the output area (no work area specified), this entry must be included. It specifies the number (r) of a general-purpose register (any one of 2 to 12), which will always contain the length of the available space remaining in the output area.

IOCS calculates the space still available in the output area, and supplies it to you in the designated register after the PUT macro is issued for a variable-length record. You then compare the length of your next variable-length record with the available space to determine if the record fits in the area. This check must be made before the record is built. If the record does not fit, issue a TRUNC macro to transfer the completed block of records to the file. Then, the present record is built at the beginning of the output area in the next block.

VERIFY=YES: This operand is included if you want to check the parity of disk records after they are written. If this operand is omitted, any records written on a disk are not verified.

WLRERR=name: This operand applies only to disk input files. It does not apply to undefined records. WLRERR specifies the symbolic name of your routine to receive control if a wrong-length record is read.

If the WLRERR operand is omitted but a wrong-length record is detected by IOCS, one of the following conditions results:

• If the ERROPT entry is included for this file, the wrong-length record is treated as an error block and handled according to your specifications for an error (IGNORE, SKIP, or name of error routine).

- If the ERROPT entry is not included, the error is ignored.

Undefined records are not checked for incorrect record length. The record is truncated when the BLKSIZE specification is exceeded.

WORKA=YES: If I/O records are processed, or built, in work areas instead of in the I/O areas, specify this operand. You must set up the work area in storage. The address of the work area, or a general-purpose register which contains the address, must be specified in each GET or PUT macro. For a GET or PUT macro, IOCS moves the record to, or from, the specified work area. WORKA=YES is required for SPNUNB and SPNBLK. When this operand is specified for a file, the IOREG operand must be omitted.

SDMODxx Macro

Sequential DASD files on FBA devices do not need SDMODxx logic module macros to be specified for them because preassembled re-entrant logic modules are loaded into the SVA at IPL time and are available to all partitions, as needed, for FBA support. These modules are functional subsets. An SDMODxx module associated with a problem program is ignored if the file is assigned to an FBA device.

Sequential DASD module generation macros differ from other IOCS module generation macros. The file characteristics are separated into ten categories, and each category has a unique macro associated with it (see Figure 3-49).

Macro	Module Generated
SDMODFI	Sequential DASD Module, Fixed-length records[1], Input file
SDMODFO	Sequential DASD Module, Fixed-length records[1], Output file
SDMODFU	Sequential DASD Module, Fixed-length records[1], Update file
SDMODVI	Sequential DASD Module, Variable-length records (including spanned records)[2], Input file
SDMODVO	Sequential DASD Module, Variable-length records (including spanned records)[2], Output file
SDMODVU	Sequential DASD Module, Variable-length records (including spanned records)[2], Update file
SDMODUI	Sequential DASD Module, Undefined records[3], Input file
SDMODUO	Sequential DASD Module, Undefined records[3], Output file
SDMODUU	Sequential DASD Module, Undefined records[3] Update file
SDMODW	Sequential DASD Module, Work file[4]

[1] RECFORM=FIXUNB or FIXBLK in DTFSD
[2] RECFORM=VARUNB, VARBLK, SPNUNB, or SPNBLK in DTFSD
[3] RECFORM=UNDEF in DTFSD
[4] RECFORM=FIXUNB or UNDEF in DTFSD

Figure 3-49. SDMOD macros.

The macro operation code and the keyword operands define the characteristics of the module. Modules for a specific file can thus be generated more quickly than if there were only one macro. A module name may be contained in the name field of the macro. The macro operation code is contained in the operation field (SDMODFI, for example). The operands are contained in the operand field. The operands for the ten macros are summarized in Figure 3-50 and explained below.

The control interval format logic lodules assume the value YES for all the following operands except RECFORM (where SPNBLK is assumed) and RPS.

CONTROL=YES: This operand is specified if a CNTRL macro is issued for the file. This entry applies to all SDMODxx macros. The module also processes any DTF in which the CONTROL parameter is not specified.

ERREXT=YES: Include this operand if non-data-transfer errors are returned to an ERROPT routine in your program or if the ERET macro is used with the DTF and module. If ERREXT is specified, ERROPT must also be specified.

ERROPT=YES: This operand applies to all SDMODxx macros. It is included if the module handles any of the error options for an error block. Logic is generated to handle any of the three options (IGNORE, SKIP, or name) regardless of which option is specified in the DTF. The module also processes any DTF in which the ERROPT operand is not specified.

If this operand is not included, your program is canceled whenever any unrecoverable error except a wrong-length record error (which LIOCS ignores) is encountered.

FEOVD=YES: This operand is specified if the forced end of volume for disk feature is desired. It allows the program to force an end of volume condition before physical end of volume occurs. When the FEOVD macro is issued, the current volume is closed, and I/O processing continues on the next volume.

HOLD=YES: This operand applies only to update files (SDMODFU, SDMODVU, and SDMODUU) and work files (SDMODW). The operand is included if the track hold function is employed. Any DTF used with the module must have the same operand.

NOTEPNT={POINTRW|YES}: This operand applies to SDMODW (work files) only. It is included if any NOTE, POINTR, POINTS, or POINTW macros are used within the module. If the operand specifies POINTRW, logic to handle only NOTE, POINTR, and POINTW is generated.

If YES is specified, the routines to handle NOTE, POINTR, POINTS, and POINTW are generated and any files that specify NOTEPNT=POINTRW in the DTF are processed.

RDONLY=YES: This operand causes a read-only module to be generated. Whenever this operand is

Operand	Required	Comments
CONTROL=YES	If the CNTRL macro is to be issued for the file.	Applies to all SDMODs.
ERREXT=YES	If the module returns non-data-transfer errors or is used with the ERET macro.	Applies to all SDMODs.
ERROPT=YES	If the module is to handle error options for an error block.	Applies to all SDMODs.
FEOVD=YES	If the FEOVD macro is to be issued for the file.	Applies to all SDMODs except SDMODW.
HOLD=YES	If the track hold function is to be employed.	Applies to update and work file logic modules.
NOTEPNT= {POINTRW|YES}	If NOTE, POINTR, POINTS, or POINTW macros are to be issued for the file.	This parameter applies to SDMODW only. The operand POINTRW generates logic for NOTE, POINTR, and POINTW. The operand YES generates logic for all macros.
RDONLY=YES	If a read-only module is to be generated.	Applies to all SDMODs.
RECFORM= {SPNUNB|SPNBLK}	If unblocked or blocked spanned records are to be processed.	Applies to SDMODVI, SDMODVO, and SDMODVU only.
RPS=SVA	If RPS support is desired.	To assemble the RPS logic modules.
SEPASMB=YES	If the module is assembled separately from the DTF.	Applies to all SDMODs.
TRUNCS=YES	If the TRUNC macro is to be issued for the file. Assumed for output files consisting of variable-length blocked records.	Applies to all SDMODs for fixed-length records.
UPDATE=YES	If SDMODW is to process the WRITE UPDATE macro.	Applies to SDMODW only.

Figure 3-50. SDMODxx operands.

specified, any DTF used with the module must have the same operand.

RECFORM= {SPNUNB|SPNBLK}: This operand is required only for SDMODVI (input files), SDMODVO (output files), and SDMODVU (update files) if RECFORM=SPNUNB or SPNBLK is specified in the DTF macro. If RECFORM is specified incorrectly, an assembler diagnostic (MNOTE) is issued, and the module generation is terminated.

RPS=SVA: This operand causes the RPS logic modules to be assembled. When this operand is used, only superset modules are generated.

SEPASMB=YES: Include this operand only if the module is assembled separately. This causes a CATALR card with the module name (standard or user-specified) to be punched ahead of the object deck and the module name to be defined as an ENTRY point in the assembly. If the operand is omitted, the assembler assumes that the module is being assembled with the DTF and the problem program and no CATALR card is punched.

TRUNCS=YES: This operand applies to all SDMOD macros for fixed-length records. It generates a logic module which can handle the TRUNC macro. This operand is assumed for VARBLK output files.

This operand is ignored if specified for VARBLK input or update files. It must be specified if any FIXBLK DASD files (processed by the module) contain short blocks embedded within them, if the input file was created with a module that specified TRUNCS, or if the DTF was specified with RECFORM=UNDEF. The module cannot process any DTF for fixed-length records in which the TRUNCS operand is not specified.

UPDATE=YES: This operand applies to the SDMODW only. It is assumed for SDMODFU, SDMODUU, and SDMODVU and generates a logic module which can handle the WRITE UPDATE macro with work files.

Standard SDMOD Names
Each name begins with a 3-character prefix (IJG) and continues with of a 5-character field corresponding to the options permitted in the generation of the module.

In SDMOD there are two module classes:

- Those which handle GET/PUT functions
- Those which handle READ/WRITE, NOTE/POINTx, and CHECK functions (work files).

Name List for GET/PUT Type Modules

SDMODxx name = IJGabcde

a = C SDMODFx specifies HOLD=YES
  = F SDMODFx does not specify HOLD=YES
  = R SDMODUx specifies HOLD=YES
  = U SDMODUx does not specify HOLD=YES
  = P SDMODVx specifies HOLD=YES and
       RECFORM=SPNBK|SPNUNB
  = Q SDMODVx does not specify HOLD=YES and speci-
       fies RECFORM=SPNBLK|SPNUNB
  = S SDMODVx specifies HOLD=YES
  = V SDMODVx does not specify HOLD=YES

b = U SDMODxU
  = I SDMODxI
  = O SDMODxO
  = W SDMODxI, RPS=SVA
  = X SDMODxO, RPS=SVA
  = Y SDITODxU, RPS=SVA

c = C ERROPT=YES and ERREXT=YES, RPS=SVA
  = E ERROPT=YES
  = Z neither is specified

d = M TRUNCS=YES and FEOVD=YES (see Note 1 be-
       low)
  = T TRUNCS=YES (see Note 1 below), RPS=SVA
  = W FEOVD=YES (see Note 2 below), RPS=SVA
  = Z neither is specified (see Note 2 below)

e = B CONTROL=YES and RDONLY=YES, RPS=SVA
  = C CONTROL=YES
  = Y RDONLY=YES
  = Z neither is specified

Notes:

1. Generated only for SDMODFx.
2. If generated for SDMODVO, TRUNC logic is available.

Name List for Workfile Type Modules (TYPEFLE=WORK)

SDMODxx name = IJGabcde

a = T HOLD=YES
  = W HOLD=YES not specified

b = C ERROPT=YES and ERREXT=YES
  = E ERROPT=YES
  = Z neither is specified

c = N NOTEPNT=YES
  = R NOTEPNT=POINTRW
  = Z NOTEPNT is not specified

d = C CONTROL=YES
  = Z CONTROL=YES is not specified

e = T RDONLY=YES and UPDATE=YES
  = U UPDATE=YES
  = Y RDONLY=YES
  = Z neither is specified

Subset/Superset SDMOD Names

Figure 3-51 illustrates the subsetting and superset-
ting allowed for SDMOD names. For the GET/PUT
type modules, four parameters allow supersetting.
For example, in the GET/PUT type module, the mo-
dule IJGFUETC is a superset of a module with the
name of IJGFUZTZ.

```
For GET/PUT Type Modules:
            +  *  +  +  +
      I  J  G  C  U  C  M  B
               F  I  E  T  Y
               +  O  Z  +  +
               R  W     W  C
               U  X     Z  Z
               +  Y
               P
               Q
               V
               +
               P
               S
               V


For Workfile Type Modules:
            +  +  +  +  +
      I  J  G  T  C  N  C  T
               W  E  R  Z  Y
               Z  Z     +
                        U
                        Z
```

+ Subsetting/supersetting permitted.
* No subsetting/supersetting permitted.
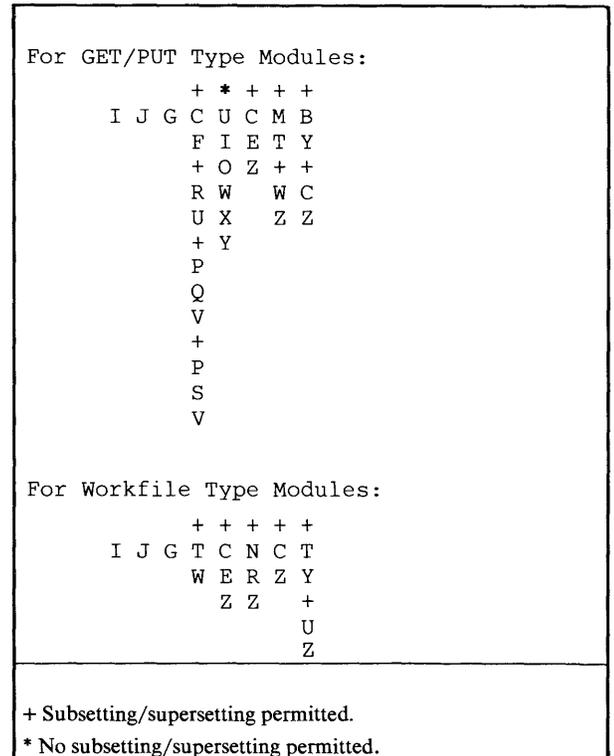
Figure 3-51. Subsetting and supersetting of SDMOD names.

CCB Macro

Name	Operation	Operand
blockname	CCB	SYSnnn,command-list-name [,X'nnnn'][,senseaddress]

A CCB (command control block) macro must be specified in your program for each I/O device controlled by physical IOCS macros. The CCB (see Figure 4-1) is necessary to communicate information to physical IOCS so that it can perform desired operations (for example, notifying your program of printer channel 9). The CCB also receives status information after an operation and makes this available to your program. You should ensure proper boundary alignment of the CCB if this is necessary for your program.

Note: In some applications, it may be preferable to use an IORB (I/O Request Block) in place of a CCB. Do this by specifying either an IORB or GENIORB macro.

blockname: The CCB macro must be given a symbolic name (blockname). This name can be used as the operand in the EXCP and WAIT macros which refer to the CCB.

SYSnnn: This operand specifies the symbolic unit for the actual I/O unit with which this CCB is associated. The actual I/O unit can be assigned to the symbolic unit by an ASSGN job control statement.

command-list-name: This operand specifies the symbolic name of the first CCW used with a CCB. This name must be the same as the name specified in the assembler CCW statement that constructs the CCW.

X'nnnn': A hexadecimal value used to set the CCB user option bits. Column 5 of Figure 4-2 gives the value used to set a user option bit 'on'. If more than one bit must be set, the sum of the values is used.

senseaddress: This operand, when supplied, indicates user error recovery (see Figure 4-2, byte 2, bit 7) and generates a CCW for reading sense information as the last field of the CCB. The name field (sense address) of the area that you supply must have a length attribute assigned of at least one byte. Physical IOCS uses this length attribute in the CCW to determine the number of bytes of sense information you desire.

CCB Format

From the above specifications, the macro sets up an area of either 16 bytes or 24 bytes. For the layout of this area and its contents see Figure 4-1.

Bytes	Residual count ①	Transmission information ②	CSW status ③	Type code ④
0				
8	reserved	CCW address ⑤	reserved	CCW address in CSW ⑥
16	optional sense CCW ⑦			
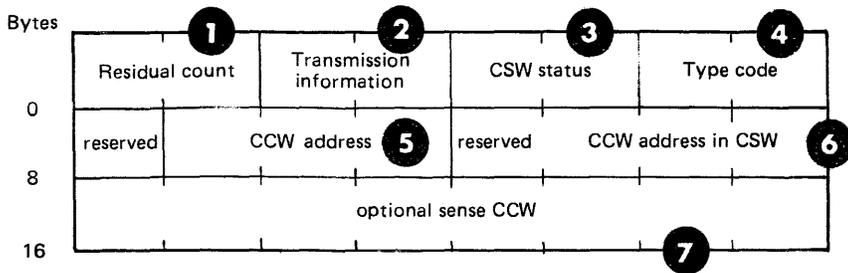
1 After a record has been transferred, IOCS places the residual count from the CSW into these two bytes. By subtracting the residual count from the original count in the CCW, your program can determine the length of the transferred record. The field is set to zero for negative values.

2 Used for transmission of information between physical IOCS and your program. For detailed information on the use and purpose of the individual bits in this field, see Figure 4-2. Your program can test any of the bits in this field using the mask given in the last column of Figure 4-2. Your program may test more than one bit by the hexadecimal sum of the test values.

All bits are set to 0 when your program is assembled unless the X'nnnn' operand is sepcified. If this operand is specified, it is assembled into these two bytes. When your program is being executed, each bit may be set to 1 by your program (to request certain functions or specific feedback information) or by physical IOCS (as a result of having detected a particular condition). Any bits that can be turned on by physical IOCS during program execution are reset to zero by PIOCS the next time an EXCP macro is executed against the same CCB.

3 Byte 4 is set to X'00' whenever an EXCP macro is issued against the CCB. For non-teleprocessing devices, a program-controlled interruption (PCI) is ignored.

The meaning of the bits in these two bytes is as follows:

Byte 4:	Byte 5:
0 = attention	0 = program-controlled interruption
1 = status modifier	1 = incorrect length
2 = control unit end	2 = program check
3 = busy	3 = protection check
4 = channel end	4 = channel data check
5 = device end	5 = channel control check
6 = unit check	6 = interface control check
7 = unit exception	7 = chaining check

If bit 5 of CCB byte 2 is set to 1 and device end results as a separate interrupt, device end will be posted.

Figure 4-1. Layout and contents of Command Control Block (CCB).

4 Contents of byte 6:

X'0u' = original CCB
X'2u' = translated CCB
X'4u' = BTAM-ES request against original CCB
X'6u' = BTAM-ES request against translated CCB
X'8u' = user-translated CCB in virtual partition

Note: if u = 0: the address in byte 7 refers to a system logical unit.
if u = 1: the address in byte 7 referes to a programmer logical unit

Contents of byte 7:
Hexadecimal representation of SYSnnn as follows:

SYSRDR = 00	SYS000 = 00
SYSIPT = 01	SYS001 = 01
SYSPCH = 02	SYS002 = 02
SYSLST = 03	•
SYSLOG = 04	•
SYSLNK = 05	•
SYSRES = 06	SYS240=F0
SYSSLB = 07	
SYSRLB = 08	
SYSUSE = 09	
SYSREC = 0A	
SYSCLB = 0B	
SYSVIS = 0C	
SYSCAT = 0D	

5 Address of CCW (or of the first of a chain of CCWs) associated with the CCB:
This is a real address if CCB byte 6 = X'2u', X'6u', or X'8u'.
This is a virtual address if CCB byte 6 = X'0u' or X'4u'.

6 Either of the following:
The CCW address contained in the CSW at channel-end interrupt for the I/O operation involving the CCB; or the address of the associated channel appendage routine if CCB byte 12 contains X'40'.

7 Bytes 16 to 23 are provided only if the sense operand was specified in the CCB macro. They accommodate the CCW for returning sense information to your program.

Byte	Bit	Condition Indicated		On Values for Third Operand in CCB Macro	Mask for Test Under Mask Instruction
		1 (ON)	0 (OFF)		
2	0 Traffic Bit (WAIT)	I/O Completed. Normally set at Channel End. Set at Device End if bit 5 is on.	I/O requested and not completed.		X'80'
	1 End of File on System Input	/* or /& on SYSRDR or SYSIPT. Byte 4, Unit exception Bit is also on.			X'40'
	3211 UCB Parity Check (line complete)[8]	Yes	No		
	2 Unrecoverable I/O Error	I/O error passed back due to program option or operator option.	No program or operator option error was passed back.		X'20'
	3[1] Accept Unrecoverable I/O Error (Bit 2 is ON)	Return to user after physical IOCS attempts to correct I/O error.[2]	Operator Option: Dependent on the Error	X'1000'	X'10'
	4[1] Return: DASD data checks, 3540 data checks,	Operator Options: Ignore, Retry, or Cancel.	Operator Options: Retry or Cancel.	X'0800'	X'08'
	2671 data checks, 1017/1018 data checks.	Ignore or Cancel.	Cancel.		
	5424/5425 not ready.	Return to user.			
	Indicate action type messages for DOC				
	5[1] Post at Device End. Specify this bit to be set on for a 2560 or 5424/5425.	Device End condition is posted: that is, byte 2, bit 0 and byte 3, bits 2 and 6 set at Device End. Also byte 4, bit 5 is set.	Device End conditions are not posted. Traffic bit is set at Channel End.	X'0400'	X'04'
	6[1] Return: Uncorrectable tape read data check (2400 series or 3420); 1018, 2560 data check, 2520 or 2540 punch equipment check; 2560, 5424/5425 read, punch, print data, and print clutch equipment checks; 3881 equipment check; 3504, 3505, or 3525 permanent errors; DASD read or read verify data check; 3211 passback requested; 3895 error codes requested. (Data checks on count not retained)[8]	Return to user; after physical IOCS attempts to correct 3211[8], tape, or DASD error; when 1018 or 2560 data check[4]; when 2560 or 5424/5425 equipment check; when 3504, 3505, 3525 permanent error (byte 3, bit 3 is also on).[4,8]	Operator Option: Ignore or Cancel for tapes, paper tape punch (1018), card punches other than 2560 and 5424/5425. Retry or Cancel for DASD, 2560, or 5424/5425.	X'0200'	X'02'
	7[1] User Error Routine	User handles error recovery.[3]	A physical IOCS error routine is used unless the CCB sense address operand is specified. The latter requires user error recovery.	X'0100	X'01'

Figure 4-2. Conditions indicated by CCB bytes 2 and 3 (Part 1 of 3).

| Byte | Bit | Condition Indicated | | On Values for Third Operand in CCB Macro | Mask for Test Under Mask Instruction |
		1 (ON)	0 (OFF)		
3	0 Data check in DASD count field. Permanent error for 3330, 3340, or 3350.	Yes-Byte 3, bit 3 is off; Byte 2, bit 2 is on.	No		X'80'
	Data check - 1287 or 1288.	Yes	No		
	MICR - SCU not operational.	Yes	No		
	3211 Print Check (equipment check).[7,8]	Yes	No		
	3540 special record transferred.	Yes	No		
	1 DASD Track overrun.	Yes	No		X'40'
	1017 broken tape.	Yes	No		
	Keyboard correction 1287 in Journal Tape Mode	Yes	No		
	3211 print quality error (equipment check)[8].	Yes	No		
	MICR intervention required.	Yes	No		
	2 End of DASD Cylinder.	Yes	No		X'20'
	Hopper Empty 1287/1288 Document Mode.	Yes	No		
	MICR - 1255/1259/1270/1275/1419, disengage.	Document feeding stopped.	No		
	1275/1419D, I/O error in external interrupt routine.	Channel data check or Bus-out check.			
	3211/2245 line position error.[5,8]	Yes	No		
	3 Tape read data check (2400 series); 2520, 2540 or 3881 equipment check; any DASD data check.	Operation was unsuccessful. Byte 2, bit 2 is also on. Byte 3, bit 0 is off.	No		X'10'
	1017, 1018 data check.	Yes	No		
	1287, 1288 equipment check.	Yes	No		
	2560, 3203, 5203, 5424/5425 read, punch, print data, and print clutch equipment checks.	Byte 2, bit 6 is also on.	No		
	3504, 3505, 3525 permanent errors.	Byte 2, bit 6 is also on.	No		
	3211 data check/print check.[8]	Yes	No		
	3540 data check.	Yes	No		
	4 Nonrecovery Questionable Condition.	Card: unusual command sequence. For DASD, no record found. 1287, 1288 document jam or torn tape. 3211 UCB parity check (command retry). 5424/5425 not ready.			X'08'
	5[1] No record found condition (retry on 2311, 2314, 2319, 3330, 3340, or 3350).	Retry command if no record found condition occurs (disk).	Set the nonrecovery questionable condition bit on and return to user.	X'0004'	X'04'

Figure 4-2. Conditions indicated by CCB bytes 2 and 3 (Part 2 of 3).

| Byte | Bit | Condition Indicated | | On Values for Third Operand in CCB Macro | Mask for Test Under Mask Instruction |
		1 (ON)	0 (OFF)		
3	6 Verify error for DASD or Carriage Channel 9 overflow.	Yes. (Set on when Channel 9 is reached only if Byte 2, bit 5 is on).	No		X'02'
	1287 document mode: late stacker select.	Yes	No		
	1288 End-of-Page (EOP).	Yes	No		
	7[1] Command Chain Retry. Specify this bit to be set on if command chaining is used for a 2560 or 5424/5425.	Retry begins at last CCW executed.[6]	Retry begins at first CCW or channel program.	X'0001'	X'01'

Notes:

1 User Option Bits. Set in CCB macro. Physical IOCS sets the other bits off at EXCP time and on when the specified condition occurs.

2 I/O program check, command reject, or tape equipment check always terminates the program.

3 You may not handle Channel Control Checks and Interface Control Checks. The occurence of a channel data check, unit check, or channel chaining check cause byte 2, bit X'20' of the CCB to turn on, and completion of posting and dequeuing to occur. I/O program and protection checks always cause program termination. Incorrect length and unit exception are treated as normal conditions (posted with completion). Also, you must request device end posting (CCB byte 2, bit X'04') in order to obtain errors after channel end.

4 Error correction feature for 1018 is not supported by physical IOCS. When a 1018 data check occurs and CCB byte 2, bit X'02' is on, control returns directly to you with CCB byte 3, bit X'10' turned on.

5 A line position error on the 3211 can occur as a result of an equipment check, data check, or FCB parity check.

6 If an error occurs, physical IOCS updates the CCW address in bytes 9 through 11 of the CCB that is used for the pertinent I/O operation and is queued to the channel queue.

7 A deleted or bad spot record has been read on a 3540 diskette. CCW chain broken, after CCW reads special record.

8 3211 remarks apply also to 3211-compatible printers (that is, with device type code of PRT1). 3895 error codes are returned in CCB byte 8. Refer to the 3895 Document Reader/Inscriber manuals for information on these error codes.

Figure 4-2. Conditions indicated by CCB bytes 2 and 3 (Part 3 of 3).

CHECK Macro

Name	Operation	Operand
[name]	CHECK	{filename|(1)} [,control-address|,(0)]

The CHECK macro prevents processing until data transfer on an I/O operation is complete. It must be issued either after a READ or WRITE macro is issued to a work file, or after a READ is issued to a MICR file.

Because of differences in the way that IOCS posts CCB transmission information bits in the DTFs, you should always issue a CHECK macro to ensure that data transfer is complete before testing these bits. If the data transfer is completed without an error or other exceptional condition, CHECK returns control to the next sequential instruction. If an error condition is encountered, control is transferred to the ERROPT address. If ERROPT is not specified, processing continues at the next instruction. If end-of-file is encountered, control transfers to the EOFADDR address.

filename|(1): The operand specifies the name of the file associated with the record to be checked or, if register notation is used, the register containing a pointer to the field that contains this name. This name is the same as that specified for the DTFXX header entry for the file.

Issuing a CHECK macro after a READ on a MICR device allows you to query the MICR document buffer (see Figure 4-3) and to specify the control-address operand:

control-address|(0): indicates the address to which control passes when a buffer is waiting for data or when the file is closed. If register notation is used, the specified register must point to a field that contains this address.

The CHECK macro determines whether the MICR document buffer contains data ready for processing, is waiting for data, contains a special nondata status, or the file (filename) is closed. If the buffer has data ready for processing, control passes to the next sequential instruction. If the buffer is waiting for data, or the file is closed, control passes to the address specified for control address, if present. If the buffer contains a special nondata status, control passes to the ERROPT routine for you to examine the posted error conditions before determining your action. (See byte 0, bits 2, 3, and 4, of the document buffer).

Return from the ERROPT routine to the next sequential instruction via a branch on register 14, or to the control address in register 0.

If the buffer is waiting for data, or if the file is closed, and the control address is not present, control is given to you at your ERROPT address specified in the DTFMR macro.

If an error, a closed file, or a waiting condition occurs (with no control-address specified) and no ERROPT address is present, control is given to you at the next sequential instruction.

If the waiting condition occurred, byte 0, bit 5 of the buffer is set to 1. If the file was closed, byte 0, bits 5 and 6 of the buffer are set to 1.

MICR Document Buffer

Buffer Status Indicators		
Byte	**Bit**	**Comment**
0	0	The document is ready for processing (you need never test this bit).
	1	Unrecoverable stacker select error, but all document data is present. You may continue to issue GETs and READs.
	2	Unrecoverable I/O error. An operator I/O error message is issued. The file is inoperative and must be closed.
	3	Unit Exception. You requested disengage and all follow-up documents are processed. The LITE macro may be issued, and the next GET or READ engages the device for continued reading.
	4	Intervention required or disengage failure. This buffer contains no data. The next GET or READ continues normal processing. This indicator allows your program to give the operator information necessary to select pockets for documents not properly selected and to determine unread documents.
	5	The program issued a READ, no document is ready for processing, byte 0, bits 0 to 2 are off, or the file is closed (byte 0, bit 6 is on). The CHECK macro interrogates this bit. **Note:** You must test bits 1 through 4 and take appropriate action. Any data from a buffer should not be processed if bits 2, 3, or 4 are on.
	6	The program has issued a GET or READ and the file is closed. Bit 5 is also on.
	7	Reserved.
1	0	Your stacker selection routine turns this bit on to indicate that batch numbering update (1419 only) is to be performed in conjunction with the stacker selection for this document. The document is imprinted with the updated batch number unless a late stacker selection occurs (byte 3, bit 2).
	1-7	Reserved. **Note:** If bits 6 or 7 (byte 2) are on, bit 0 is ignored by the external interrupt routine. With the 1419 (dual address) only, batch numbering update cannot be performed with the stacker selection of auto-selected documents.
2*	0	For 1419 or 1275 (dual address) only. An auto-select condition occurred after the termination of a READ but before a stacker select command. The document is auto-selected into the reject pocket.
	1-3	Reserved.
	4	Data check occured while reading. You should interrogate byte 3 to determine the error fields.
	5	Overrun occurred while reading. Byte 3 should be interrogated to determine the error fields. Overruns cause short length data fields. When the 1419 or 1275 is enabled for fixed-length data fields, bit 4 is set.
	6-7	The specific meanings of bits 6 and 7 depend on the device type, the model, and the Engineering Change level of the MICR reader; but if either bit is on, the document(s) concerned is (are) auto-selected into the reject pocket. 1. **1412 or 1270:** Bit 6 on indicates that a late read condition occurred. Bit 7 on indicates that a document spacing error occurred. (Unique to the 1270: both the current document and the previous document are auto-selected into the reject pocket when this bit is on. This previous document reject cannot be detected by IOCS, and byte 5 of its document buffer does not reflect that the reject pocket was selected). 2. **1275 and 1419 (single address) without engineering change #125358:** Bit 6 indicates that either a late read condition or a document spacing error occurred. Bit 7 indicates a document spacing error for the current document. 3. **1255, 1259, 1275, and 1419 (single or dual address) with engineering change #125358:** Bit 6 indicates that an auto-select condition occurred while reading a document. The bit is set at the termination of the READ command before the stacker select routine receives control. Bit 7 is always zero.

Figure 4-3. MICR document buffer format (Part 1 of 2)

Buffer Status Indicator (Continued)		
Byte	Bit	Comment
3*	0	Field 6 valid.**
	1	Field 7 valid.**
	2	A late stacker selection (unit check late stacker select on the stacker select command). The document is auto-selected into the reject pocket.
	3	Amount field valid (or field 1 valid).**
	4	Process control field valid (or field 2 valid).**
	5	Account number field valid (or field 3 valid).**
	6	Transit field valid (or field 4 valid).**
	7	Serial number field valid (or field 5 valid).**
		Notes:
		1. For the 1270, bits 3-7 are set to zero when the fields are read without error.
		2. For the 1255, 1259, 1275, and 1419, bits 3-7 are set on when each respective field, including bracket symbols, is read without error. This applies to bits 0, 1, and 3-7 on the 1259 and 1419 model 32.
		3. For the 1255, 1259, 1275, and 1419, unread fields contain zero bits. Errors are indicated when an overrun or data check condition occurs while reading the data field.

* Byte 2 (bits 4, 5, 6, and 7) and byte 3 contain MICR sense information.
** Only for the 1259 model 34 or 1419 model 32. Bits 0 and 1 are not used for other models.

Byte	Bit	Comment
4		Inserted pocket code determination by your stacker select routine. Whenever byte 0, bits 2, 3, or 4 are on, this byte is X'00' because no document was read and your stacker selection routine was not entered. Whenever auto-selection occurs, this value is ignored. A no-op (X'03') is issued to the device, and a reject pocket value (X'CF') is placed in byte 5. The pocket codes are (byte 2, bit 6 or 7 on):
		Pocket A - X'AF'* Pocket 5 - X'5F'
		Pocket B - X'BF'** Pocket 6 - X'6F' Except 1270
		Pocket 0 - X'0F' Pocket 7 - X'7F' models 1 and 3
		Pocket 1 - X'1F' Pocket 8 - X'8F'
		Pocket 2 - X'2F' Pocket 9 - X'9F'
		Pocket 3 - X'3F' Reject
		Pocket 4 - X'4F' Pocket - X'CF'
5		The actual pocket selected for the document. The contents are normally the same as that in byte 4.
		Note:
		1. X'CF' is inserted whenever auto-selection occurs (byte 2, bit 6; byte 2, bit 7; byte 2, bit 0; or byte 3, bit 2). These conditions may result from late READ commands, errant document spacing, or late stacker selection.
		a. Start I/O for stacker selection is unsuccessful (byte 0, bit 1).
		b. An I/O error occurs (for example, invalid pocket code) on the 1419 (dual address) secondary control unit when selecting this document.

Additional User Work Areas
This additional buffer area can be used as a work area and/or output area. Its size is determined by the DTFMR ADDAREA operand. The only size restriction is that this area, plus the 6-byte status indicators and data portion must not exceed 256 bytes. This area may be omitted.

Document Data Area
The document data area immediately follows your work area. The data is right-adjusted in the document data area. The length of this data area is determined by the DTFMR RECSIZE operand.

* 1275, 1419, and 1270 models 2 and 4 only.
** 1275 and 1419 only.

Figure 4-3. MICR document buffer format (Part 2 of 2)

CLOSE and CLOSER Macros

The CLOSE or CLOSER macro is used to deactivate previously opened files; they end the association between a logical file declared in a program and a specific physical file on an I/O device.

A file may generally be closed at any time, with the following exceptions;
- Console files need not be closed; the CLOSE(R) macro is invalid for files defined by means of the DTFCN.
- Files assigned to an FBA device may not be closed in an ERROPT routine.

Files (such as on an FBA device) that use control interval format must be closed in order to ensure that data in the control interval buffer be physically written on the FBA device.

No further commands can be issued to the closed file until it is reopened.

The format of the CLOSE macro is

Name	Operation	Operand
[name]	{CLOSE| CLOSER}	{filename1|(r1)} [,filename2|(r2)]...

The format of the CLOSER macro is the same except that you code CLOSER instead of CLOSE in the operation field.

When CLOSER is specified, the symbolic address constants that CLOSER generates from the parameter list are self-relocating. When CLOSE is specified, the symbolic address constants are not self-relocating.

To write the most efficient code in a multiprogramming environment it is recommended that CLOSER be used.

Enter the symbolic name of the file (assigned in the DTF header entry) in the operand field. A maximum of 16 files may be closed by one macro by entering additional filename parameters as operands. Alternatively, you can load the address of the filename in a register and specify the register using ordinary register notation. The high-order 8 bits of this register must be zeros. For CLOSER, the address of filename may be preloaded into any of the registers 2 through 15. For CLOSE, the address of filename may be preloaded into register 0 or any of the registers 2 through 15.

Notes:

1. If you use register notation, we recommend that you follow the practice of using only registers 2 through 12.

2. If CLOSE or CLOSER is issued to an unopened tape input file, the option specified in the DTF rewind option is performed. If CLOSE or CLOSER is issued to an unopened tape output file, no tapemark or labels are written.

CNTRL Macro

Name	Operation	Operand
[name]	CNTRL	{filename|(1)} ,code[,n1][,n2]

The CNTRL (control) macro provides commands for magnetic tape units, card devices, printers, DASDs, and optical readers. (Note that if the device is an FBA DASD, the CNTRL macro is treated as a no-op or null operation.) Commands apply to physical non-data operations of a unit and are specific to the unit involved. They specify such functions as rewinding tape, card stacker selection, and line spacing on a printer. For optical readers, commands specify marking error lines, correcting a line for journal tapes, document stacker selecting, or ejecting and incrementing documents. The CNTRL macro does not wait for completion of the command before returning control to you, except when certain mnemonic codes are specified for optical readers.

CNTRL usually requires two or three parameters. The first parameter must be the name of the file specified in the DTF header entry. It can be specified as a symbol or in register notation.

The second parameter is the mnemonic code for the command to be performed. This must be one of a set of predetermined codes (see Figure 4-4).

The third parameter, n1, is required whenever a number is needed for stacker selection, immediate printer carriage control, or for line or page marking on the 3886. The fourth parameter, n2, applies to delayed spacing or skipping or to timing mark check on the 3886. In the case of a printer file, the parameters n1 and n2 may be required.

The CNTRL macro must not be used for printer or punch files if the data records contain control characters and the entry CTLCHR is included in the file definition.

Whenever CNTRL is issued in your program, the DTF CONTROL operand must be included (except for DTFMT and DTFDR) and CTLCHR must be omitted. If control characters are used when CONTROL is specified, the control characters are ignored and treated as data.

IBM Unit	Mnemonic Code	n_1	n_2	Command
2400, 3410, 3420 Series Magnetic Tape Units	REW			Rewind Tape
	RUN			Rewind and Unload Tape
	ERG			Erase Gap (Writes Blank Tape)
	WTM			Write Tapemark
	BSR			Backspace to Interrecord Gap
	BSF			Backspace to Tapemark
	BSL			Backspace Logical Record
	FSR			Forward Space to Interrecord Gap
	FSF			Forward Space to Tapemark
	FSL			Forward Space Logical Record
1442, 2520 Card Read Punch	SS	1 2		Select Stacker 1 or 2
	E			Eject to Stacker 1 (1442 only)
2540 Card Read Punch 3504, 3505 Card Readers 3525 Card Punch	PS	1 2 3		Select Stacker 1, 2, or 3 (For 3504, 3505, and 3525, 3 Defaults to Stacker 2)
2560 Multi-Function Card Machine	SS	1 2 3 4 5		Select Stacker 1, 2, 3, 4, or 5
2596 Card Read Punch	SS	1 2		Select Stacker 1 for Read, or Stacker 3 for Punch / Select Stacker 2 for Read, or Stacker 4 for Punch
5424/5425 Multi-Function Card Unit	SS	1 2 3 4		Select Stacker 1, 2, 3, or 4
1403, 1443, 3203, PRT1, 3800, 5203 Printers 3525 Card Punch with Print Feature [1]		See Note		
	SP	c	d	Carriage Space 1, 2, or 3 lines
	SK	c	d	Skip to Channel c and/or d (For 3525, a Skip to Channel 1 is Valid Only for Print Only Files)
1403, 5203 Printers with Universal Character Set Feature or 3203, PRT1, or 3800 Printers [1]	UCS	ON OFF		Data Checks are Processed with an Operator Indication / Date Checks are ignored and Blanks are Printed
PRT1 Printer [1]	FOLD			Print Upper Case Characters for any Byte with Equivalent Bits 2–7
	UNFOLD			Print Character Equivalents of any EBCDIC Byte
2311, 2314, 2319, 3330, 3333, 3340, 3344 DASD [2]	SEEK			Seek to Address
3881 Optical Mark Reader	PS	1 2		Select Stacker 1 or 2
1287 Optical Reader	MARK			Mark Error Line in Tape Mode
	READKB			Read 1287 Keyboard in Tape Mode
	EJD			Eject Document
	SSD	1 2 3 4		Select Stacker A, B, Reject, or Alternate Stacking Mode
	ESD	1-4		Eject Document and Select Stacker
	INC			Increment Document at Read Station
1288 Optical Page Reader	ESD	1 3		Select Stacker A / Reject Stacker (R)
	INC			Increment Document at Read Station
3886 Optical Character Reader	DMK	name (r) number		Page mark the document when it is stacker selected as specified in parameter n_1.
	LMK	name (r) number, number		Line mark the document when it is stacker selected as specified in parameter n_1.
	ESP	1 2	name (r) number	Eject and stacker select the current document to stacker A or B. Perform line mark station timing mark check as indicated in parameter n_2.

c = An Integer that Indicates Immediate Printer Control (before printing).

d = An Integer that Indicates a Delayed Printer Control.

[1] Note: PRT1 refers to 3211-compatible printers (that is, with a device type of PRT1).

[2] Note: This includes the 3350 operating in 3330 compatibility mode.

Figure 4-4. CNTRL macro command codes

DISEN Macro

Name	Operation	Operand
[name]	DISEN	{filename|(1)}

This macro stops the feeding of documents through
the magnetic character reader or optical
reader/sorter. The program proceeds to the next
sequential instruction without waiting for the disen-
gagement to complete. You should continue to issue
GETs or READs until the unit exception bit (byte 0,
bit 3), of the buffer status indicators is set on (see
Figure 4-3).

The only required operand specifies the name of
the file to be disengaged. This name is the same as
that specified for the DTFMR header entry for the
file. The operand can be specified either as a symbol
or in register notation.

DSPLY Macro

Name	Operation	Operand
[name]	DSPLY	{filename|(1)},(r2),(r3)

The DSPLY macro displays the document field on
the 1287 display scope. A complete field may be
keyboard-entered if a 1287 read error makes this
type of correction necessary. An unreadable charac-
ter may be replaced by the reject character either by
the operator (if processing in the on-line correction
mode) or by the device (if processing in the off-line
correction mode). You may then use the DSPLY
macro to display the field in error.

DSPLY always requires three parameters. The first
parameter is the symbolic name specified in the
DTFOR header entry for the 1287 file. The second
parameter specifies a general-purpose register (any
from 2 to 12) into which the problem program
places the address of the load format CCW giving the
document coordinates for the field to be displayed.
When the DSPLY macro is used in the COREXIT rou-
tine, the address of the load format CCW can be ob-
tained by subtracting 8 from the 3-byte address that
is right-justified in the fullword location beginning
at filename+32. (The high-order fourth byte of this
full word should be ignored.) If the DSPLY macro is
not used in the COREXIT routine, you must deter-
mine the load format CCW address. The third par-
ameter specifies a general-purpose register (2
through 12) into which you place the address of the
load format CCW giving the coordinates of the refer-
ence mark associated with the displayed field.

ENDFL Macro

Name	Operation	Operand
[name]	ENDFL	{filename|(0)}

The ENDFL (end file load mode) macro ends the
mode initiated by the SETFL macro. The name of
the file to be loaded is the only parameter required,
and is the same as the name specified in the DTFIS
header entry for the file. The filename can be speci-
fied either as a symbol or in register notation. Regis-
ter notation is necessary if your program is to be
self-relocating. The ENDFL macro must be issued
only after a SETFL and before a CLOSE or CLOSER.

The ENDFL macro performs an operation similar
to CLOSE or CLOSER for a blocked file. It writes the
last block of data records, if necessary, and then
writes an end-of-file record after the last data re-
cord. Also, it writes any index entries that are need-
ed followed by dummy index entries for the unused
portion of the prime data extent.

ERET Macro

Name	Operation	Operand
[name]	ERET	{SKIP|IGNORE|RETRY}

This macro enables your program's ERROPT or
WLRERR routine to return to IOCS and specify an
action to be taken. The macro applies only to DTFIS,
DTFMT, DTFSD and DTFDU files with the ERREXT
operand specified.

The SKIP operand passes control back to the logic
module to skip the block of records or control inter-
val in error and process the next one. For disk or
diskette output, an ERET SKIP is treated as an ERET
IGNORE.

The IGNORE operand passes control back to the
module to ignore the error and continue processing.

The RETRY operand causes the module to retry
the operation that resulted in the error. With
MTMOD for any error or with SDMOD wrong-length
record errors, RETRY cancels the job with an invalid
SVC message.

ESETL Macro

Name	Operation	Operand
[name]	ESETL	{filename|(1)}

The ESETL (end set limit) macro ends the sequential
mode initiated by the SETL macro. For filename

specify the same name as was specified in the DTFIS header entry. The name can be specified as a symbol or in register notation. If the records are blocked, ESETL writes the last block back if a PUT was issued. Register notation is necessary if your program is to be self-relocating.

Note: If ADDRTR and/or RANSEQ are specified in the same DTF, ESETL should be issued before issuing a READ or WRITE; another SETL can be issued to restart sequential retrieval. Sequential processing must always be terminated by issuing an ESETL macro.

EXCP Macro

Name	Operation	Operand
[name]	EXCP	{blockname|(1)} [,REAL]

The EXCP (execute channel program) macro requests physical IOCS to start an input/output operation for a particular I/O device.

Physical IOCS determines the device from the CCB or IORB control block specified by blockname. Physical IOCS places the block in a queue of such blocks and returns control to the problem program. Physical IOCS causes the channel program to be executed as soon as the channel and device are available. I/O interruptions are used to process I/O completion and to start I/O for requests if the channel or device was busy at the time the EXCP was executed.

blockname: is the virtual address of the control block established for the device. It can be given as a symbol or in register notation.

REAL: indicates that the addresses in the CCWs and the address in the control block pointing to the first CCW have already been translated into real addresses; the operand causes the DOS/VSE CCW translation routine to be skipped. (For a program running in real mode, the operand is ignored.)

In your program, the EXCP macro with the REAL operand must be preceded by the PFIX macro that causes DOS/VSE (1) to page in those program pages which contain the pertinent control block, channel program, I/O areas, and IDA (indirect address) words (if used) and (2) to fix these pages in their page frames.

Notes:

1. If the I/O area being used crosses page boundaries, the data address in the appropriate CCW(s) must point to the required indirect data address words within your program; in addition, bit 37 (the IDA bit) of these CCWs must be set to 1.

2. A channel program has to start with:

   - a long seek command in the case of a CKD DASD.

   - a define extent command in the case of an FBA DASD.

   The data chaining and, in S/370 mode, also the IDA bit must be set to zero.

FEOV Macro

Name	Operation	Operand
[name]	FEOV	{filename|(1)}

The FEOV (force end-of-volume) macro is used for files on magnetic tape (programmer logical units only) to force an end-of-volume condition before sensing a reflective marker. This indicates that processing of records on one volume is considered finished, but that more records for the same logical file are to be read from, or written on, a following volume. For system units, see "SEOV Macro".

The name of the file is the only parameter required. The name can be specified either as a symbol or in register notation.

When physical IOCS macros are used and DTFPH is specified for standard label processing, FEOV may be issued for output files only. In this case, FEOV writes a tapemark, the standard trailer label, and any user-standard trailer labels if DTFPH LABADDR is specified. When the new volume is mounted and ready for writing, IOCS writes the standard header label and user-standard labels, if any.

FEOVD Macro

Name	Operation	Operand
[name]	FEOVD	{filename|(1)}

The FEOVD (force end-of-volume for disk) macro is used for either input or output files to force an end-of-volume condition before it actually occurs. This indicates that processing of records on one volume is finished, but that more records for the same logical file are to be read from, or written on, the following volume. If extents are not available on the new volume, or if the format-1 label is posted as the last volume of the file, control is passed to the EOF address specified in the DTF.

The name of the file is the only required operand. The name can be specified either symbolically or in register notation.

GENIORB Macro

Name	Operation	Operand
[name]	GENIORB	[ADDRESS={name1|(S,name1)|(1)}] [,LENGTH=fieldlength] ,CCW={name2|(S,name2)|(r2)} ,{DEVICE=SYSxxx| LOGUNIT={name3|(S,name3)| (r3)}} [,FIXLIST={name4|(S,name4)| (r4)}] [,FIXFLAG=(option11,...)] [,IOFLAG=(option21,...)] [,ERREXIT={name5|(S,name5)| (r5)}] [,ECB={name6|(S,name6)|(r6)}]

The GENIORB macro generates an IORB
(Input/Output Request Block). The block is gener-
ated at the time of program execution. For the lay-
out and contents of an IORB, see Figure 4-5. The
IORB is an alternative to the CCB; instead of speci-
fying a CCB in the EXCP macro, the address of an
IORB is given. If used in a program that is to be exe-
cuted under DOS/VSE operating in ECPS:VSE mode,
the IORB allows for the specification of areas to be
page-fixed for the I/O operation. Such areas include
the IORB and the channel programs themselves and
all input/output areas. Specifying those areas frees
the DOS/VSE page-fixing routines from having to
scan the channel programs to determine which areas
are to be fixed.

The GENIORB does not provide for SENSE CCWs.
GENIORB and CCB may be used interchangeably,
depending on the required functions.

In S/370 mode, the GENIORB macro is accepted,
but a fixlist, if specified, is not used.

After execution of the macro,
register 1 contains the address of the IORB, and
register 15 contains the return code from an
implicit GETVIS.

For a detailed display in your assembly, showing
the IORB fields and their meaning, issue the IORB
macro, with the (only) operand DSECT=YES.

ADDRESS= {name1|(S,name1)|(1)}: If speci-
fied, this operand gives the name of the area in
which the IORB is to be generated. The ADDRESS
operand can be specified only together with the
LENGTH operand.

Omitting the ADDRESS operand indicates that the
required area is to be obtained thru an implicit
GETVIS issued by DOS/VSE.

LENGTH=field length: This operand gives the
length of the field provided for IORB generation.
The value must be given as a selfdefining term. If
this operand is omitted a default value equal to the

length of the IORB will be used; however, the assem-
bler issues an MNOTE. If the ADDRESS operand is
omitted, LENGTH will not be used.

CCW={name2|(S,name2)|(r2)}: This operand
gives the name of the first CCW used with the IORB.
The name must be the same as the name specified in
the assembler CCW statement that builds the CCW.

DEVICE=SYSxxx: This operand specifies the
symbolic unit for the actual I/O unit with which the
IORB is associated.

LOGUNIT= {name3|(S,name3)|(r3)}: This ope-
rand describes the device in logical unit format. It
points to a halfword with the same format as a logi-
cal unit number (bytes 6 and 7) in a CCB; see Figure
4-1 provided in context with the discussion of the
CCB macro.

For a description of the FIXLIST, FIXFLAG, and
IOFLAG operands, refer to the IORB macro.

ERREXIT= {name5|(S,name5)|(r5)}: ERREXIT is
the address of a routine to be executed should
DOS/VSE be unable to obtain the required virtual
storage. If the ERREXIT operand is omitted, failure
to obtain virtual storage causes DOS/VSE to cancel
the program (task).

ECB= {name6|(S,name6)|(r6)}: This operand
specifies the address of the ECB to be posted when
I/O is complete. For a more detailed description of
the ECB operand, refer to the IORB macro.

GET Macro

Name	Operation	Operand
[name]	GET	{filename|(1)} [,workname|,(0)]

GET makes the next sequential logical record from
an input file available for processing in either an
input area or in a specified work area. It is used for
any input file in the system, and for any type of re-
cord: blocked or unblocked, fixed or variable
length, and undefined.

If GET is used with a file containing checkpoint
records, the checkpoint records are bypassed auto-
matically.

filename: This operand is required. The parame-
ter value must be the same as specified in the header
entry of the DTF macro for the file from which the
record is to be retrieved. The filename can be speci-

fied as a symbol or in either special or ordinary register notation. The latter is necessary to make your programs self-relocating.

workname: This is an optional parameter specifying the work area name or a register (in either special or ordinary register notation) containing the address of the work area. The work area address should never be preloaded into register 1. This parameter is used if records are to be processed in a work area which you define yourself (for example, using a DS instruction). If the operand is specified, all GETs for the named file must use a register or a workname. Using the second operand causes GET to move each individual record from the input area to a work area. The workname parameter is not valid for the 3881. You also cannot specify the WORKA operand in the DTFCD for the 3881.

In conjunction with optical reader input, this macro can be used only to retrieve records from a journal tape on a 1287.

IORB Macro

Name	Operation	Operand
[name]	IORB	DSECT=YES or CCW=name1[,DEVICE =SYSxxx, [,FIXLIST=name2] [,FIXFLAG=(option11,...)] [,IOFLAG=(option21,...)] [,ECB=name3]

The IORB macro generates an IORB (Input/Output Request Block). The block is generated when your program is being assembled. For the layout and contents of an IORB, see Figure 4-5.

The IORB is an alternative to the CCB: instead of specifying a CCB in the EXCP macro, the address of an IORB is given. If used in a program that is to be executed under DOS/VSE operating in ECPS:VSE mode, the IORB macro allows for the specification of areas to be page-fixed for the I/O operation. Such areas include the IORB and the channel programs themselves and all input/output areas. Specifying those areas frees the DOS/VSE page-fixing routines from having to scan the channel programs to determine which areas are to be fixed.

The IORB does not provide for SENSE CCWs. IORB and CCB may be used interchangeably, depending on the required functions.

In S/370-mode the IORB is accepted, but the specified fixlist (see the discussion of operand FIXLIST=name2, below) is not used.

CCW=name1: This operand give the name of the first CCW used with the IORB. This name must be the same as the name specified in the assembler CCW statement that builds the CCW.

DEVICE=SYSxxx: This operand specifies the symbolic unit for the actual I/O unit with which this IORB is associated.

FIXLIST=name2: This operand specifies the address of the first of two or more fixlist parts or of the only fixlist part. In ECPS:VSE mode, the FIXLIST operand is required unless FIXFLAG=FIXED is specified. Each fixlist part consists of one or more 8-byte entries plus an end or chaining indicator as shown in Figure 4-6.
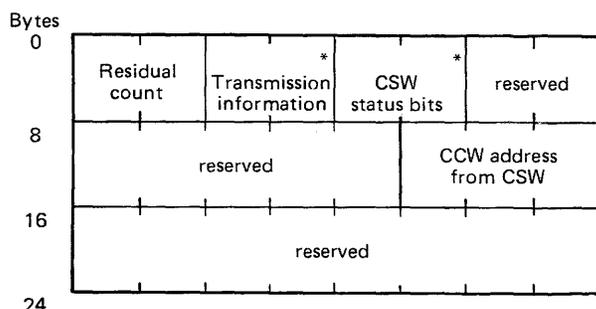


* Same as for a CCB (see Figure 4-1)

Figure 4-5. Layout and contents of the I/O Request Block (IORB)
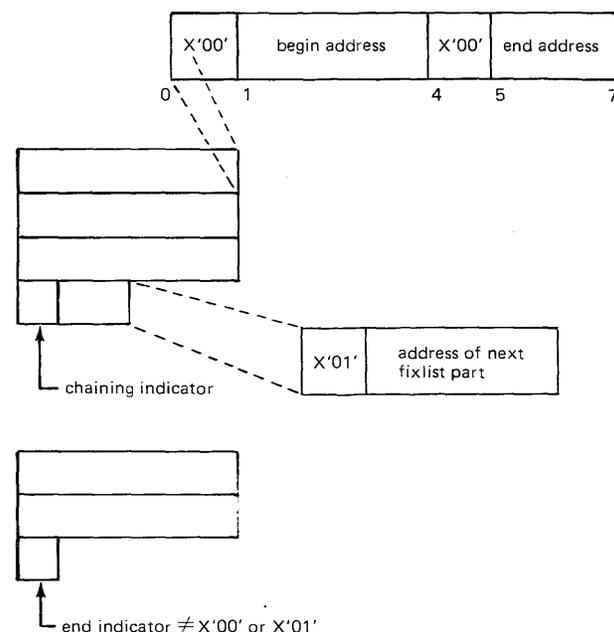
Figure 4-6. Layout of fixlist

In a fixlist entry, begin address and end address are the addresses of the first and the last byte of an area which has to be fixed for the I/O request (begin address ≤ end address; entries with begin address > end address will be ignored). Each entry describes a storage area that is accessed by the channel during the I/O request; that is, an area containing the channel program, an input/output area, or the IORB.

Duplicate entries and entries describing overlapping storage areas are allowed. As a result, certain areas may be covered more than once by the fixlist. The system will compress the fixlist so that each page to be fixed for the channel program is covered only once. However, specifying FIXFLAG=(COMPRESSED) indicates that this service is not desired.

FIXFLAG=(option11,...): This operand specifies a list of options which apply to the I/O fixing procedure.

The options you can specify are:

COMPRESSED   to indicate that the system does not have to compress the fixlist. Use this option if the fixlist is already compressed, that is: each page to be fixed for the I/O request is covered only once by the fixlist.

FIXED   to indicate that all areas which should be fixed for the I/O operation have already been fixed by the user, there is no fixlist. In S/370 mode, this specification is ignored.

IOFLAG=(option21,...): A list of options may be specified which apply to I/O interrupt handling:

POSTDE   to indicate that device end is to be posted.

POSTERR   to indicate that an unrecoverable I/O error is to be accepted.

SKIPERP   to indicate that error recovery by the system is to be skipped.

ECB=name3: This operand specifies the address of the ECB to be posted when I/O is complete. The traffic bit (byte 2, bit 0) of the ECB must have been cleared before issuing the EXCP macro. The ECB operand must be specified if a fixlist is used.

Note: If FIXFLAG=FIXED is specified, the ECB must have been PFIXed.

DSECT=YES: If the operand is specified, it should be the only one. Any other parameters specified in the macro are ignored and an appropriate MNOTE is generated by the assembler.

Specifying DSECT=YES causes the assembler to display, as a DSECT structure, the IORB and the meaning of its fields.

LBRET Macro

Name	Operation	Operand
[name]	LBRET	{1|2|3}

The LBRET macro is issued in your subroutines when you have completed processing labels and wish to return control to IOCS. LBRET applies to subroutines that write or check DASD or magnetic tape user-standard labels, write or check tape non-standard labels, or check DASD extents. The operand used – 1, 2, or 3 – depends on the function to be performed. The functions and operands are explained below.

Checking User Standard DASD Labels: IOCS passes the labels to you one at a time until the maximum allowable number is read (and updated), or until you signify you want no more. In the label routine, use LBRET 3 if you want IOCS to update (rewrite) the label just read and pass you the next label. Use LBRET 2 if you simply want IOCS to read and pass the next label. If an end-of-file record is read when LBRET 2 or LBRET 3 is used, label checking is automatically ended. If you want to eliminate the checking of one or more remaining labels, use LBRET 1.

Writing User Standard DASD Labels: Build the labels one at a time and use LBRET to return to IOCS, which writes the labels. Use LBRET 2 if you want control returned to you after IOCS writes the label. If, however, IOCS determines that the maximum number of labels has already been written, label processing is terminated. Use LBRET 1 if you wish to stop writing labels before the maximum number of labels is written.

Checking User Standard Tape Labels: IOCS reads and passes the labels to you one at a time until a tapemark is read, or until you indicate that you do not want any more labels. Use LBRET 2 if you want to process the next label. If IOCS reads a tapemark, label processing is automatically terminated. Use LBRET 1 if you want to bypass any remaining labels.

Writing User Standard Tape Labels: Build the labels one at a time and return to IOCS, which writes the labels. When LBRET 2 is used, IOCS returns control to you (at the address specified in LABADDR)

after writing the label. Use LBRET 1 to terminate the label set.

Writing or Checking Nonstandard Tape Labels:
You must process all your nonstandard labels at once. Use LBRET 2 after all label processing is completed and you want to return control to IOCS.

LITE Macro

Name	Operation	Operand
[name]	LITE	{filename|(1)} [,light-switches|,(0)]

This macro lights any combination of pocket lights on a 1419 magnetic character reader or 1275 optical reader/sorter. Before using the LITE macro, the DISEN macro must be issued to disengage the device. Processing of the documents should be continued until the unit exception bit (byte 0, bit 3) of the document buffer status indicators is set on (see Figure 4-3). When this bit is on, the follow-up documents have been processed, the MICR reader has been disengaged, and the pocket LITE macro can be issued.

The first operand is the name of the file; this name is the same as that specified for the DTFMR header entry for the file. The second operand indicates a 2-byte area containing the pocket light switches. Both operands can be given either as a symbol or in register notation.

The bit configuration for the pocket light switch area is shown in Figure 4-7. The pocket lights that are turned on should have their indicator bits set to 1. If an error occurs during the execution of the pocket lighting I/O commands, bit 7 in byte 1 is set to 1. This error condition normally indicates that the pocket light operation was unsuccessful.

NOTE Macro

Name	Operation	Operand
[name]	NOTE	{filename|(1)}

The NOTE macro obtains identification for a physical record or logical block that is read or written during processing. At least one READ or WRITE operation should be successfully completed by means of the CHECK macro before issuing the NOTE macro. To NOTE a desired record successfully, the POINTR, POINTS, or POINTW macros must not be issued between CHECK and NOTE.

For magnetic tape, the last record read or written in the specified file is identified by the number of physical records read or written from the load point. The physical record number is returned in binary in the three low-order bytes of register 1. The high-order byte contains binary zero.

For CKD DASD, the binary number returned in register 1 is in the form cchr, where

cc = cylinder number,
h = track number,
r = record number within the track.

Register 0 contains the unused space remaining on the track following the end of the identified record.

For FBA devices, register 1 contains an address relative to the begining of the file in the form cccb, where ccc is the relative number of the current control interval (origin 0), and b is the relative block number within the current CI (origin 1). Register 0 contains the length of the longest logical block that could completely fit in the CI following the NOTEd logical block.

You must provide a four- or six-byte field and store in it the record identification and the remaining capacity so that it can be used later by a POINTR or POINTW macro to find the NOTEd record again. The two-byte track or CI capacity remaining is needed only when a WRITE SQ is to follow the POINTR or POINTW.

OPEN and OPENR Macros

The OPEN or OPENR macro activates all files.

When OPENR is specified, the symbolic address constants that OPENR generates from the parameter list are self-relocating. When OPEN is specified, the symbolic address constants are not self-relocating. The format of the OPEN macro is

Name	Operation	Operand
[name]	OPEN	{filename1|(r1)} [,filename2|,(r2)] ...

Bits	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
Pocket Lights	A	B	0	1	2	3	4	5	6	7	8	9	Reserved			Error indicator bit

Figure 4-7. Bit configuration for pocket light switch area of the 1419

The format of the OPENR macro is the same except that you code OPENR instead of OPEN in the operation field.

Enter the symbolic name of the file (DTF filename) to be opened in the operand field. A maximum of 16 files may be opened with one OPEN or OPENR by entering the filenames as additional operands.

Self-relocating programs using LIOCS must use OPENR to activate all files, including console files. In addition to activating files for processing, OPENR relocates all address constants within the DTF tables (zero constants are relocated only when they constitute module address).

To write the most efficient code in a multiprogramming environment it is recommended that OPENR be used.

POINTR Macro

Name	Operation	Operand
[name]	POINTR	{filename|(1)} ,{address|(0)}

The POINTR macro repositions the file specified by filename to the record identified by previouly issuing a NOTE macro. The filename may be expressed either as a symbol or in register notation.

address| (0): specifies a virtual storage location in which is stored either a four-byte record identifier or a four-byte record identifier plus a two-byte track or CI capacity. The four- or six-byte number must be in the form obtained from the NOTE macro. The two-byte track or CI capacity is required only when a WRITE SQ is to be issued following the POINTR.

If a READ follows the POINTR, the NOTEd record is the record read.

For magnetic tape, a WRITE must not follow a POINTR.

For DASD, if a WRITE UPDATE follows the POINTR, the NOTEd record is written (or overwritten). If a WRITE SQ follows the POINTR, the record after the NOTEd record is written (or overwritten) and, on CKD DASD, the remainder of the track is erased (overwritten with zeros). On FBA devices, the remainder of the CI is erased (overwritten with zeros) and an SEOF is written (the following CI is also overwritten with zeros).

POINTS Macro

The POINTS macro repositions a file to its beginning.

Name	Operation	Operand
[name]	POINTS	{filename|(1)}

For a tape file, the tape is rewound. If the file contains any header labels, they are bypassed, and the tape is positioned to the first record following the label set.

For disk, the file is repositioned to the lower limit of the first extent. A POINTS should not be followed by a WRITE UPDATE. If a POINTS is followed by a WRITE SQ, the first record in the file is overwritten. For CKD DASD, the remainder of the track is then erased (overwritten with zeros). For FBA devices, the remainder of the CI is erased (overwritten with zeros) and an SEOF is written (the following CI is also overwritten with zeros).

POINTW Macro

Name	Operation	Operand
[name]	POINTW	{filename|(1)} ,{address|(0)}

The POINTW macro repositions the file specified by filename to the record following the record identified by previously issuing a NOTE macro. The filename may be expressed either as a symbol or in register notation.

address| (0): specifies a virtual storage location in which is stored either a four-byte record identifier or a four-byte record identifier plus a two-byte track or CI capacity. The four- or six-byte number must be in the form obtained from the NOTE macro. The two-byte track or CI capacity is required only when a WRITE SQ is to be issued following the POINTW.

If a READ follows the POINTW, the record after the NOTEd record is read.

For DASD, if a WRITE UPDATE follows the POINTW, the record *after* the NOTEd record is written (or overwritten). If a WRITE SQ follows the POINTW, the record *after* the NOTEd record is written (or overwritten) and, on CKD DASD, the remainder of the track is erased (overwritten with zeros). On FBA devices, the remainder of the CI is erased (overwritten with zeros) and an SEOF is written (the following CI is also overwritten with zeros).

PRTOV Macro

Name	Operation	Operand
[name]	PRTOV	{filename|(1)}, {9|12} [,routine-name|,(0)]

The PRTOV (printer overflow) macro is used with a printer file to specify the operation to be performed when a carriage overflow condition occurs. To use this macro, the PRINTOV=YES operand must be included in the DTFPR or DTFSR.

PRTOV requires either two or three parameters. The first parameter must be the filename, written either as a symbol or in register notation. The second parameter must specify the number of the carriage control channel (9 or 12) used to indicate the overflow. When an overflow condition occurs, IOCS restores the printer carriage to the first printing line on the form (channel 1), and normal printing continues.

The third parameter is required if you prefer to branch to your own routine on an overflow condition, rather than skipping directly to channel 1. It specifies the name of the routine, as a symbol or in register notation. However, the name should never be preloaded into register 1.

If you specify the third parameter, IOCS does not restore the carriage to channel 1.

PUT Macro

Name	Operation	Operand
[name]	PUT	{filename|(1)} [,workname|,(0)] ,STLSP= {controlfield|(r)} ,STLSK= {controlfield|(r)}

PUT writes, prints, or punches logical records which are built directly in the output area or in a specified work area. PUT can be used for any sequential output file defined by a DTF macro, and for any type of record: blocked or unblocked, fixed or variable length, and undefined. It operates much the same as GET but in reverse. It is issued after a record has been built.

filename: This operand is required. The parameter value must be the same as specified in the header entry of the DTF for the file being built. The operand can be specified as a symbol or in either special or ordinary register notation. Use register notation if your program is to be self-relocating.

workname: An optional parameter specifying the work area name or a register (in either special or ordinary register notation) containing the address of the work area. The work area address should never be preloaded into register 1. This parameter is used if records are built in a work area which you define yourself (for example, using a DS instruction). If the operand is specified, all PUTs to the named file must use a register or a workname. Using the second operand causes PUT to move each record from the work area to the output area.

Individual records for a logical file may be built in the same work area or in different work areas. Each PUT macro specifies the work area where the completed record was built. However, only one work area can be specified in any one PUT macro.

Whenever a PUT macro transfers an output data record from an output area (or work area) to an I/O device, the data remains in the area until it is either cleared or replaced by other data. IOCS does not clear the area. Therefore, if you plan to build another record whose data does not use every position of the output area or work area, you must clear that area before you build the record. If this is not done, the new record will contain interspersed characters from the preceding record.

STLSP=control field: This optional operand specifies a control byte that allows for spacing while using the selective tape listing feature on the 1403 printer. To use this feature, the operand STLST=YES must be specified in the DTFPR. Up to 8 paper tapes may be independently spaced. The control byte is set up like any other data byte in virtual storage. You can also use ordinary register notation to provide the address of the control byte. Registers 2 through 12 are available without restriction. You determine the spacing (which occurs after printing) by setting on the bits corresponding to the tapes you want to space. The correspondence between control byte bits and tapes is as follows:

Control byte bits	0	1	2	3	4	5	6	7
Tape position	8	7	6	5	4	3	2	1

The tape position 1 is the leftmost tape on the selective tape listing device.

Note: Double-width tapes must be controlled by both bits of the control field.

STLSK=control field: This optional operand specifies a control byte that allows for skipping while using the selective tape listing feature on the 1403 printer. To use this feature, the operand STLIST=YES must be specified in the DTFPR. Up to 8

paper tapes may be independently skipped. The control byte is set up like any other data byte in virtual storage. You can also use ordinary register notation to provide the address of the control byte. Registers 2 through 12 are available without restriction. You determine the skipping (which occurs after printing) by setting on the bits corresponding to the tapes you want to skip. The correspondence between control byte bits and tapes is shown in the figure under "STLSP=control field", above.

PUTR Macro

Name	Operation	Operand
[name]	PUTR	{filename|(1)} [,{workname1|(0)}, {workname2|(2)}]

The PUTR (PUT with reply) macro is used for the display operator console, to issue a message to the operator which requires operator action and which will not be deleted from the display screen until the operator has issued a reply.

You may also use PUTR with the 3210 or 3215 console printer-keyboard, in which case PUTR functions the same as PUT followed by GET for these devices, but provides the message non-deletion code for the display operator console. Use of PUTR for the 3210 or 3215 is therefore recommended for compatibility if your program may at some time be run on the display operator console instead of the 3210 or 3215.

Use PUTR for fixed unblocked records (messages). Issue PUTR after a record has been built.

filename: This operand is required. The parameter value must be the same as specified in the header entry of the DTFCN for the file being built. The filename can be specified as a symbol or in either special or ordinary register notation. The latter is necessary to make your programs self-relocating.

workname1|(0): An optional parameter specifying the output work area name or a register (in either special or ordinary register notation) containing the address of the output work area. The work area address should never be preloaded into registers 1 or 2. This parameter is used if records are built in a work area which you define yourself (for example, using a DS instruction). The length of the work area is defined by the BLKSIZE parameter of the DTFCN macro. If workname1 is specified, workname2 must also be specified.

workname2|(2): An optional parameter specifying the input work area name or a register (in either special or ordinary register notation) containing the address of the input work area. The work area address should never be preloaded into registers 0 or 1. This parameter is used if records are built in a work area which you define yourself (for example, using a DS instruction). The length of the work area is defined by the INPSIZE parameter of the DTFCN macro. If workname2 is specified, workname1 must also be specified.

RDLNE Macro

Name	Operation	Operand
[name]	RDLNE	{filename|(1)}

The RDLNE macro provides selective on-line correction when processing journal tapes on the 1287 optical reader. This macro reads a line in the on-line correction mode while processing in the off-line correction mode. RDLNE should be used in the COREXIT routine only, or else the line following the one in error will be read in on-line correction mode.

If the 1287 cannot read a character, IOCS first resets the input area to binary zeros and then retries the line containing the character that could not be read. If the read is unsuccessful, you are informed of this condition via your error correction routine (specified in DTFOR COREXIT). The RDLNE macro may then be issued to cause another attempt to read the line. If the character in the line still cannot be read, the character is displayed on the 1287 display scope. The operator keys in the correct character, if possible. If the operator cannot readily identify the defective character, he may enter the reject character in the error line. This condition is posted in filename+80 for your examination. Wrong-length records and incomplete read conditions are also posted in filename+80.

This macro requires only one parameter, the symbolic name of the 1287 file from which the record is to be retrieved. This name is the same as that specified in the DTFOR header entry for the file.

READ Macro

Name	Operation	Operand
[name]	READ	{filename| (1)} {,SQ, {area| (0)} [,length|,(r)|S]| ,KEY| ,OR, {name| (r)} | ,ID| ,DR, {name| (r) |number,number} | ,MR}

The READ macro transfers a record or part of a record from an input file to an area in virtual storage.

filename\| (1): Specifies the name of the file from which the record is to be read. The name is the same as that specified in the DTF header entry.

SQ: Required for sequential files.

area\| (0): The name of the input area used by a sequential file.

length\| (r) \| S: Used only for sequential files of undefined format (RECFORM=UNDEF). Specifies the actual number of bytes to be read, or the register where the number is to be found. S specifies that the entire record is to be read.

KEY: For ISAM, KEY is required. For DAM, specifies that the record reference is to be by record key (control information in the key area of the DASD record).

OR: Signifies that the file is for a 1287 or 1288 optical character reader.

name\| (r): Specifies the CCW list address to be used to read a document from the 1287 or 1288.

ID: For DAM, specifies that the reference is to be by ID (identifier in the count area of the record).

DR: Indicates a 3886 Optical Character Reader is the input device.

The third parameter specifies the line number to be read and the format record for the line in one of three ways:

- **name\| (r)** provides the symbolic address of a 2-byte hexadecimal field containing the line number in the first byte and the format record number in the second byte.
- **(r)** provides the number of the register that contains the address of the two-byte hexadecimal field.

- **number,number** provides the decimal line number to be read (any number from 1 through 33), followed by the format record number used to reaad the line (0-63).

MR: signifies that the file is for a magnetic ink character reader (MICR).

RELSE Macro

Name	Operation	Operand
[name]	RELSE	{filename|(1)}

The RELSE (release) macro is used with blocked input records read from a DASD device, or with blocked spanned records read from, or updated on, a DASD device. This macro is also used with blocked input records read from magnetic tape.

The macro allows you to skip the remaining records in a block and continue processing with the first record of the next block when the next GET macro is issued. When used with blocked spanned records, RELSE makes the next GET skip to the first segment of the next record.

The symbolic name of the file, specified in the DTF header entry, is the only parameter required for this macro. It can be specified as a symbol or in register notation.

RESCN Macro

Name	Operation	Operand
[name]	RESCN	{filename|(1)} ,(r1),(r2),[n1][,n2]

The RESCN macro selectively rereads a field on a document if one or more defective characters make this type of operation necessary. The field is always right-justified into the area (normally within IOAREA1) that was originally intended for this field as specified in the CCW. The macro first resets this area to binary zeros.

Note: For the 1287 models 3 and 4 and the 1288, this macro can only be used with READ BACKWARD commands. If used with READ FORWARD commands, the input area is not cleared. When 1288 unformatted fields are read, the RESCN macro should not be used.

The first parameter, filename, specifies the symbolic name of the 1287D file as specified in the DTFOR header entry for the file.

The second parameter, r1, specifies a general-purpose register from 2 to 12 into which the program places the address of the load format CCW.

The third parameter, r2, specifies a general-purpose register from 2 to 12 into which the program places the address of the load format CCW for reading the reference mark.

The previous three parameters are always required, and result in one attempted reread for the field.

The fourth parameter, n1, allows you to specify the number of attempts (one to nine allowed) to reread the unreadable field. If this parameter is omitted, one is assumed.

The fifth parameter, n2, indicates one more reread which forces on-line correction of any unreadable character(s) by individually projecting the unreadable character(s) on the 1287 display scope.

The operator must key in a correction (or reject) character(s). This operand cannot be used for 1288 processing.

SECTVAL Macro

Name	Operation	Operand
[name]	SECTVAL	[DDKR={name1|(0)}] [,DVCTYP=name2]

The SECTVAL macro calculates the sector value of the address of the requested record on the track of a disk storage device when RPS is used. The macro returns this value in register 0.

The sector value is calculated from data length, key length, and record number information. Values are calculated for fixed or variable length and for keyed and non-keyed records.

DDKR= {name1|(0)}: The information needed to calculate the sector value should be specified in the 4-byte field at name1, or in the specified register. If no operand is specified, register 0 is the default and should contain the necessary information. The four bytes of information have the format DDKR, where

DD= a 2-byte field which specifies:
- for fixed length records, the data length of each record, or
- for variable length records, the number of bytes used on the track, excluding standard R0 length up to the current record. Bit 0 of the first byte must be set on.

K= a 1-byte field indicating the key length:
- for fixed length records, the actual key length must be specified;
- for variable length records, any non-zero value is sufficient to indicate the presence of keys.

Note: For non-keyed records the value should be 0.

R= a 1-byte record number field which specifies the number of the record of which the sector value is being requested.

DVCTYP=name2: The device type code is specified at name2. If no operand is specified, it is assumed that byte 0 of register 1 contains the code. The following device type codes (which are the same as the device type codes generated in the DTF) are valid for:

IBM 3330, models 1 and 2: X'04'
IBM 3330, model 11: X'05'
IBM 3340: X'08', X'09', or X'0A'
IBM 3350: X'07'

The calculated sector value is returned in register 0. If any errors are detected in calculating the sector value, a no-operation sector value (X'FF') is returned.

SEOV Macro

Name	Operation	Operand
[name]	SEOV	filename

The SEOV (system end-of-volume) macro must only be used with physical IOCS to automatically switch volumes if SYSLST or SYSPCH are assigned to a tape output file. SEOV writes a tapemark, rewinds and unloads the tape, and checks for an alternate tape. If none is found, a message is issued to the operator who can mount a new tape on the same drive and continue. If an alternate unit is assigned, the macro fetches the alternate switching routine to promote the alternate unit, opens the new tape, and makes it ready for processing. When using this macro, you must check for the end-of-volume condition in the CCB.

SETDEV Macro

Name	Operation	Operand
[name]	SETDEV	{filename|(1)} ,{phasename|(r)}

The SETDEV macro changes format records during execution of the program. When the new format record has been loaded into the 3886, control returns to the next sequential instruction in your program. If the operation is not successful, the completion code is posted at EXITIND and control is passed to the COREXIT routine, or the job is canceled. If you issue the SETDEV macro and no documents remain to be processed and the end-of-file key has been

pressed on the device, control is passed to the end-of-file routine.

The first parameter, filename, specifies the same name as that used in the DTFDR header entry. Register notation must be used if your program is to be self-relocating.

The second parameter specifies the name of the format record to be loaded, phasename; or indicates the register containing the address of an 8-byte area that contains the phasename.

SETFL Macro

Name	Operation	Operand
[name]	SETFL	{filename|(0)}

The SETFL (set file load mode) macro causes ISAM to set up the file so that the load or extension function can be performed. This macro must be issued whenever the file is loaded or extended.

When loading a file, SETFL preformats the last track of each track index. When extending a file, SETFL preformats only the last track of the last track index plus each new track index for the extension of the file. This allows prime data on a shared track to be referenced even though no track indexes exist on the shared track.

The name of the file loaded is the only parameter required for this macro and is the same as that specified in the DTFIS header entry for the file. It can be specified as a symbol or in register notation. Register notation is necessary if your program is to be self-relocating.

SETL Macro

Name	Operation	Operand
[name]	SETL	{filename|(r)}
		, {id-name|(r)|KEY|BOF|
		GKEY}

The SETL (set limits) macro initiates the mode for sequential retrieval and initializes ISAM to begin retrieval at the specified starting address. The first operand (filename) specifies the same name as that used in the DTFIS header entry, as a symbol or in register notation. Register notation is necessary if your program is to be self-relocating.

The second operand specifies where processing is to begin.

If you are processing by the record ID, the operand id-name or (r) specifies the symbolic name of

the 8-byte field in which you supply the starting (or lowest) reference for ISAM use. This field contains the information shown in Figure 4-8.

If processing begins with a key you supply, the second operand is KEY. The key is supplied in the field specified by the DTFIS KEYARG operand. If the specified key is not present in the file, an indication is given at filenameC.

BOF specifies that retrieval is to start at the beginning of the logical file.

Selected groups of records within a file containing identical characters or data in the first locations of each key can be selected by specifying GKEY (generic key) as the second operand. GKEY allows processing to begin at the first record (or key) within the desired group. You must supply a key that identifies the significant (high order) bytes of the required group of keys. The remainder (or insignificant) bytes of the key must be padded with blanks, binary zeros, or bytes lower in collating sequence than any of the insignificant bytes in the first key of the group to be processed. For example, a GKEY specification of D6420000 would permit processing to begin at the first record (or key) containing D642XXXX, regardless of the characters represented by the x's. Your program must determine when the generic group is completed. Otherwise, ISAM continues through the remainder of the file.

Note: If the search key is greater than the highest key on the file, the filename status byte is set to X'10' (no record found).

The ESETL (end set limit) macro should be issued before issuing a READ or WRITE if ADDRTR and/or RANDSEQ are specified in the same DTF. Another SETL can be issued to restart sequential retrieval. Sequential processing must always be terminated by issuing an ESETL macro.

TRUNC Macro

Name	Operation	Operand
[name]	TRUNC	{filename|(1)}

The TRUNC (truncate) macro is used with blocked output records written on DASD or magnetic tape. It allows you to write a short block of records. Blocks do not include padding. Thus, the TRUNC macro can be used for a function similar to that of the RELSE macro for input records. That is, when the end of a category of records is reached, the last block can be written and the new category can be started at the beginning of a new block.

Note that if DEVICE=FBA is specified on the DTF, TRUNC will not necessarily cause a physical write to the FBA DASD unless PWRITE is also specified.

Byte	Identifier	Contents in Hexadecimal	Information
0	m	02-F5	Number of the extent in which the starting record is located
1-2	bb	0000 (disk)	Always zero for disk
3-4	cc	0000-00C7 (2311, 2314, 2319) 0000-0193 (3330, 3333) 0000-015B (3348 model 35) 0000-02B7 (3348 model 70)	Cylinder number for disk: for 2311, 2314, 2319: 0-199 for 3330, 3333: 0-403 for 3340 with 3348 model 35: 0-347 for 3340 with 3348 model 70: 0-695
5-6	hh	0000-0009 (2311) 0000-0013 (2314, 2319) 0000-0012 (3330, 3333) 0000-000B (3340)	Head position for disk
7	r	01-FF	Record location

Figure 4-8. Field supplied for SETL processing by record ID

The symbolic name of the file, specified in the DTF header entry, is the only parameter required in this macro.

WAIT Macro

Name	Operation	Operand
[name]	WAIT	{blockname|(1)}

Issue this macro whenever your program requires that an I/O operation (started by an EXCP macro) be completed before execution of the program continues. For example, transferring data (a physical record) to virtual storage must be completed before data can be added or moved to another area of virtual storage, or otherwise processed. When WAIT is executed in a batched job environment, processing is suspended until the traffic bit (byte 2, bit 0) of the related CCB or IORB is turned on. Then, processing automatically continues and the data can be processed. In a multiprogramming environment, the supervisor gives control to another program until the traffic bit is set on.

The blockname (specified as a symbol or in register notation) of the CCB or IORB established for the I/O device is the only operand required. This is also the same name as that specified in the EXCP macro for the device.

WAITF Macro

Name	Operation	Operand
[name]	WAITF	{filename1|(r1)} [,filename2|,(r2)] ...

The WAITF macro is issued to ensure that the transfer of a record is complete. It is valid for both DAM and ISAM, but for SAM only with MICR and OCR

devices. Filename is the same as that used in the DTF header entry, and may be specified either as a symbol or in register notation. Note that multiple filenames are valid only when using SAM to read MICR records.

The WAITF macro is issued after any READ or WRITE for a file and before the succeeding READ or WRITE for the same file. If the I/O operation is not completed when WAITF is issued, the active partition is placed in a wait state until the data transfer is completed. This allows processing of programs in other partitions while waiting for completion. When data transfer is complete, and if no errors were encountered, processing continues with the next sequential instruction. If an error is encountered, control passes to the error-handling routine provided for in the DTF.

If, however, you are using the multiple filename format of the WAITF macro while using MICR records, and if any of the files have records or errors ready to be processed, control remains in the partition and processing continues with the instruction following the WAITF.

WRITE Macro

Name	Operation	Operand
[name]	WRITE	{filename|(1)} {,{SQ|UPDATE},{area|(0)}| [,length|,(r)]| ,KEY| ,ID| ,AFTER[,EOF]| ,NEWKEY| ,RZERO}

The WRITE macro transfers a record from virtual storage to an output file.

filename|(1): Filename specifies the same name as that used in the DTF header entry. Register notation must be used if your program is to be self-relocating.

SQ|UPDATE: For sequential files, specify SQ for magnetic files, or for disk work files for a formatting write (count, key, and data). Specify UPDATE for a non-formatting write (data only).

area|(0): For sequential files, specifies the name, as a symbol or in register notation, of the output area used by the file.

length|(r): Specifies the actual number of bytes to be written on a sequential file. Is used only for records of undefined format (RECFORM=UNDEF).

KEY: For indexed sequential files, specify KEY for random updating. For direct access files, specify KEY to write in a location determined by the record key (control information in the key area of the records).

ID: For DA files, specify ID to write in a location determined by the record identifier in the count area of the records.

AFTER: For DA files, specify AFTER to write a record after the last record written, regardless of key or identifier.

EOF: Optional: applies only to the WRITE...AFTER form of the macro. Specify to write an end-of-file on a track after the last record on the track.

NEWKEY: For indexed sequential files only; specify NEWKEY to write a new (not updated) record in the file. When loading or extending the file, preceed the WRITE filename,NEWKEY with a SETFL macro and follow it with an ENDFL macro. When adding a record after sequential retrieval, issue an ESETL macro before writing the record.

RZERO: For DA files, specify RZERO to reset the capacity record of a track to its maximum value and erase the track after record zero.

ATTACH Macro

Name	Operation	Operand
[name]	ATTACH	{entrypoint|(S,entrypoint)|(r1)}
		,SAVE= {savearea|(S,savearea)|(r2)}
		[,ECB= {ecbname|(S,ecbname)|(r3)}]
		[,ABSAVE= {savearea|(S,savearea)|(r4)}]
		[,MFG= {area|(S,area)|(r5)}]

A subtask can be intitiated only by issuing the ATTACH macro within the main task. The part of the subtask containing the entry point must be in storage before the subtask can be successfully attached.

If register notation is used in any of the macro operands, register 0 and 1 should not be specified.

entrypoint |(S,entrypoint)|(r1): The operand specifies the entrypoint of the subtask.

SAVE=(savearea|(S, savearea)|(r2): The operand must provide the address of the save area for the subtask. The save area is 96 or 128 bytes in length depending upon whether or not the supervisor provides floating-point support (for S/370 mode, FP=YES must have been specified in the CONFG generation macro).

If an interrupt occurs while the subtask is in control, the system saves in this save area the subtask's interrupt status information, general purpose registers, and (depending on floating-point support) the floating-point registers (see Figure 5-1).

Before issuing the ATTACH macro, provide a subtask name in the first eight bytes of the save area. The name is used to identify the subtask in the event of a possible abnormal termination condition.

ECB= {ecbname|(S,ecbname)|(r3)}: The operand must be specified if other tasks can be affected by this subtask's termination, or if the ENQ and DEQ macro are used within the subtask. This parameter is the name of the subtask's event control block (ECB), and is a fullword defined within your program. The ECB may be any 4-byte (or larger) field where in byte 2, bit 0 is the termination indicator and bit 1 is the abnormal indicator. The remaining bits of the four bytes are reserved. At the time a subtask is attached, bits 0 and 1 of byte 2 are set to 0. When a subtask terminates, the supervisor sets byte 2, bit 0 of the ECB to 1. In addition, byte 2, bit 1 is set to 1 when the subtask terminates abnormally; that is, if task termination is not the result of issuing

Save area without floating-point option:
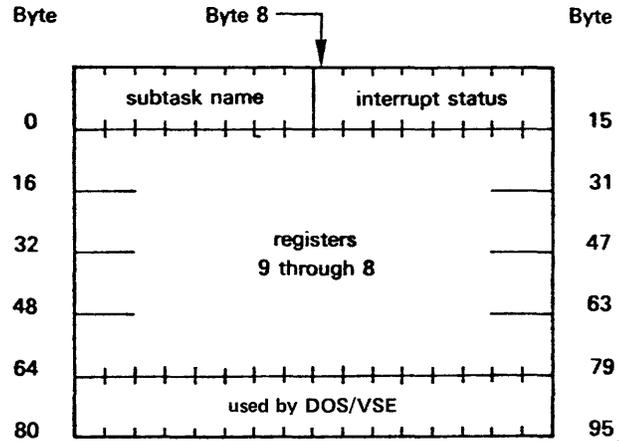

Save area with floating-point option
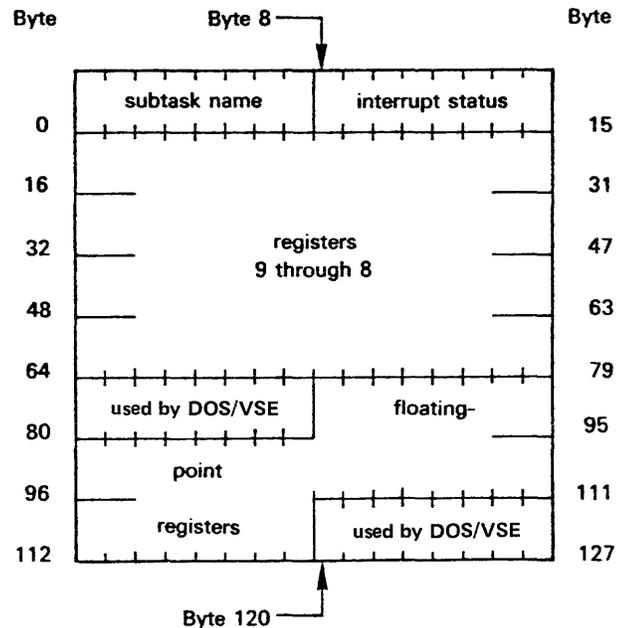
Figure 5-1. Subtask save area.

the CANCEL, DETACH, DUMP, or EOJ macros.

ABSAVE= {savearea|(S,savearea)|(r4)}: Specify this operand only if the subtask is to use the main task abnormal termination routine (see STXIT macro), that is, if it does not provide an abnormal termination routine of its own. Your program can have separate subtask STXIT AB routines with or without a main task STXIT AB routine, or it can have neither. The parameter specified in this operand must be the

address of a 72-byte (doubleword-aligned) STXIT save area for the subtask. When an abnormal termination occurs, the supervisor saved the old PSW and general registers 0 through 15 in this area before the exit is taken.

MFG= {area|(S,area)|(r5)}: The operand is required if the program which issues the ATTACH macro is to be reenterable. It specifies the address of a 64-byte dynamic storage area, that is, storage which your program obtained through a GETVIS macro. This area is required for system use during execution of the macro.

If the ATTACH macro successfully initiates a subtask, the subtask is given higher priority than the main task, and control passes to the subtask. Register 1 of the subtask contains the address of the main task save area, and the contents of the main task registers 2 through 15 remain as they were prior to issuing the ATTACH; they are also passed to the appropriate fields in the subtask save area. The address in register 1 can be used as the second operand of a POST macro later in the job if task-to-task communication is desired. Upon return from a successful ATTACH, the main task register 0 contains the address of the byte immediately following the subtask save area, as determined by the supervisor. Register 0 therefore can be tested to ascertain whether the supervisor contains the floating-point option.

The maximum possible number of subtasks is determined by subtracting the number of partitions from the number 15. For example, the maximum possible number of subtasks is 13 for a 2-partition system or 10 for a 5-partition system. In the event that the maximum possible number of subtasks is already attached, any attempt to attach another subtask will be unsuccessful. In this case, the main task will keep control and register 1 (main task) will contain the address of an ECB within the supervisor that will be posted when the system can initiate another subtask. Register 1 will also have the bit 0 on to aid the main task in testing for an unsuccessful ATTACH.

Note: If your program uses VSAM files, you should provide a STXIT AB and PC macro and issue a CLOSE or TCLOSE for the files before canceling the subtask.

CALL Macro

The format of the CALL macro is:

Name	Operation	Operand
[name]	CALL	{entrypoint|(15)}
		[,(parameterlist)]

The CALL macro passes control from one program to a specified entry point in another program.

entrypoint|(15): specifies the entry point to which control is passed. If the symbolic name of an entry point is specified, an instruction

   L 15,=V(entrypoint)

is generated as part of the macro expansion. The linkage editor makes the called program part of the calling program phase. The symbolic name must be either the name of a control section (CSECT) or an assembler language ENTRY statement operand in the called program. Control is given to the called program at this address. The called program resides in storage throughout execution of the calling program. This wastes storage if the called program is not needed throughout execution of the calling program.

If register 15 is specified, the entrypoint address must have been loaded into that register. Control is given to the called program at the address in register 15. Specifying register 15 preceded by a LOAD macro is most useful when the same program is called many times during execution of the calling program, but is not needed in storage throughout execution of the calling program.

parameterlist: specifies one or more addresses (relocatable or absolute expressions) to be passed to the called program. Terms in the address must not be indexed. The addresses must be written in a sublist, with each address separated from the next by a comma. As part of the macro expansion, a parameter list is generated. It consists of a fullword for each address. Each fullword is aligned on a fullword boundary and contains the address to be passed in its three low-order bytes. The high-order bit in the last fullword is set to 1. When the called program is entered, register 1 (the parameter list register) contains the address of the parameter list.

CANCEL Macro

Name	Operation	Operand
[name]	CANCEL	[ALL]

Issuing the CANCEL macro in a subtask abnormally terminates the subtask without branching to any abnormal termination routine. A CANCEL ALL macro issued in a subtask, or a CANCEL issued in the main task, abnormally terminates all processing in the partition (job). Job termination in multitasking causes all abnormal termination exits (via STXIT AB) to be taken for each task except the one that issued the CANCEL macro. Once these exits are taken, the job is terminated. Upon task termination, system

messages (using the first eight bytes of each subtask save area) are issued to identify each subtask terminated.

If the CANCEL macro is issued without an operand, the macro must not contain a comment unless the comment begins with a comma. If CANCEL ALL is issued, the macro may include a comment.

If the DUMP option was specified, and SYSLST is assigned, a system dump will occur

- if a CANCEL ALL macro is issued by a subtask, or
- if a CANCEL macro is issued by a main task with subtasks attached.

CDLOAD Macro

Name	Operation	Operand
[name]	CDLOAD	{phasename|(1)} [,PAGE={NO|YES}] [,RETPNF={NO|YES}]

The CDLOAD macro loads the phase specified in the first parameter from the core image library into the partition GETVIS area. The phase is loaded only if it is not yet in either the partition GETVIS area or the SVA. CDLOAD returns control to the phase which issued the macro.

The CDLOAD macro must not be used for a phase that has been linked as a member of an overlay structure. Instead, use the LOAD macro without specifying a load address.

If a phase is to be loaded, CDLOAD determines the size of the phase, acquires the appropriate amount of GETVIS storage, and loads the phase into that storage.

After successfully loading the phase or if loading is not required (because the phase is already in the partition GETVIS area or in the SVA), registers contain values as follows:

      register 0: the load address,
      register 1: the entry point,
      register 14: the length of the phase.

phasename: For phasename, specify the name of the required phase. If register notation is used, the register must contain the address of an 8-byte field that holds the phase name as an alphameric character string.

Page= {NO|YES}: If you want to have the phase loaded on a page boundary, specify PAGE=YES.

RETPNF= {NO|YES}: determines whether the

issuing phase is canceled if the phase to be loaded does not exist in the core image library. With RETPNF=YES, the phase is not canceled; instead, control is returned to the issuing phase with the appropriate return code.

Return Codes in Register 15

After execution of the macro, register 15 contains one of the following return codes:

  0 CDLOAD completed successfully.
  4 The size of the partition's GETVIS area is OK.
  8 The length of requested GETVIS storage is negative.
 12 No more storage is available in the GETVIS area.
 16 The partition CDLOAD directory (also known as anchor table) is full.
 20 The phase does not exist in the core image library (this return code occurs only with RETPNF=YES).
 32 A hardware (storage) failure occured in the requested real partition GETVIS area.

CHAP Macro

Name	Operation	Operand
[name]	CHAP	

The CHAP macro lowers the priority of the issuing subtask. This issuing subtask now becomes the subtask with the lowest priority of all the subtasks within the partition.

A CHAP macro issued by the main task is ignored.

The supervisor must have been generated with multitasking support, otherwise the task issuing the CHAP macro will be canceled.

CHKPT Macro

Name	Operation	Operand
[name]	CHKPT	SYSnn, {restart-address|(r1)} [,end-address|,(r2)] [,tpointer|,(r3)] [,dpointer|,(r4)] [,filename|,(r5)]

The CHKPT macro is used to record the status of your program so that the program, should its execution be terminated before it has completed processing, may be restarted using job control. The partition in which the program is to be restarted must start at the same location as when the program was checkpointed, and its end address must not be lower than the end address at checkpoint time. If the

CHKPT macro is successfully executed, control is returned with the checkpoint number in unpacked decimal format in register 0. If it is unsuccessful and the checkpoint has not been taken, register 0 contains zero and the reason is printed on SYSLOG.

Note: If a program using routines in the SVA is being checkpointed, you must make sure that SVA routines occupy the same locations on restart, should a restart become necessary.

Special register notation cannot be used with any of the CHKPT macro operands.

All VSAM files should be closed before the CHKPT macro is issued.

SYSnnn: Specifies the logical unit on which the checkpoint information is to be stored. It must be an EBCDIC magnetic tape or a disk pack.

restart address|(r1): Specifies a symbolic name of the instruction (or register containing the address) at which execution is to restart if processing must be continued later.

end address|(r2): A symbolic name assigned to (or register containing the address of) the uppermost byte of the program area required for restart. This address must be higher than the highest address of storage occupied by any phase loaded into the partition. The address should be a multiple of 2K. If the address is not a multiple of 2K, it is rounded to the next 2K boundary. If this operand is omitted, all storage allocated to the partition (other than the GETVIS area) is checkpointed.

The specified end address is ignored if any GETVIS was executed in the partition. (Note that GETVIS storage may have been requested by included IBM routines). In this case again, all storage allocated to the partition is checkpointed.

tpointer|(r3): Address of an 8-byte field containing 2 V-type address contants used in repositioning magnetic tape at restart time. The address may be a symbolic address or contained in a register. For details, see the section "Repositioning Magnetic Tape" in the *DOS/VSE Macro User's Guide*.

dpointer|(r4): Address of a DASD operator verification table, used to allow the operator to verify DASD volume serial numbers at restart time. May be a symbolic address or contained in a register.

filename|(r5): The name of the associated DTFPH; used only for checkpoint records on disk.

COMRG Macro

Name	Operation	Operand
[name]	COMRG	[REG = r]

The COMRG macro places the address of the communication region of the partition from which the macro is issued into the specified register. If the operand is omitted, register 1 is assumed.

CPCLOSE Macro

Name	Operation	Operand
[name]	CPCLOSE	[arglist|(r1)]

In spooling programs written in Basic Assembler Language, the CPCLOSE macro can be used to issue a CP CLOSE command to VM/370 in order to release a print or punch file for output.

Note: The CPCLOSE macro is valid only with a DOS/VSE system with VSE/Advanced Functions installed and a supervisor generated with the VM=YES option specified on the SUPVR macro.

arglist|(r1): This operand specifies a 16-byte argument list whose format is described below and which must be set up before issuing the macro. If the argument list name is specified, DOS/VSE loads the address into register 1. If a register is specified, it is assumed to contain the address of the argument list and this address is loaded into register 1. If no operand is specified, register 1 is assumed to contain the address of the argument list.

0	Packed device address	EBCDIC device address	Job name
0	2	4	8 15

Packed device address = unit record device address of the device to be closed (in packed format).
EBCDIC device address = unit record device address of the device to be closed (in EBCDIC format).
Job name = name of the job.

Return codes in Register 15

X'00'–Successful completion of CPCLOSE macro.

X'04'–Device is invalid, no CLOSE is issued.

X'08'–VM=YES support not included in the supervisor.

DEQ Macro

Name	Operation	Operand
[name]	DEQ	{rcbname|(0)}

A task releases a resource by issuing the DEQ macro. If other tasks are enqueued on the same RCB, the

DEQ macro frees from their wait condition all other tasks that were waiting for that resource. In such cases, the highest priority task either obtains or maintains control. A task that attempts to dequeue a resource that was not enqueued or that was enqueued by another task is abnormally terminated. Dequeuing under these two conditions within an abnormal termination routine results in a null operation instruction.

rcbname|(0): The operand is the same as that in the ENQ macro and specifies the address of the RCB.

DETACH Macro

Name	Operation	Operand
[name]	DETACH	[SAVE={savearea|(1)}]

The DETACH macro terminates execution of a task. A subtask is normally terminated by issuing a DETACH macro, and no operand is required in this case. The main task can also terminate a subtask it initiated by issuing the DETACH macro with an operand. The operand provides the address of the save area specified in the ATTACH macro for the subtask to be terminated. If the main task issues the DETACH macro without specifying an operand, all programs in the partition are terminated abnormally. The DETACH macro sets byte 2, bit 0 of the ECB to 1 (if specified in the ATTACH macro) to indicate normal termination. All tasks waiting on this ECB are taken out of the wait state, and the highest priority task obtains control.

Note: If your program uses VSAM files, ensure that these files are closed before you issue this macro.

DTL Macro

Name	Operation	Operand
[name]	DLT	[NAME=resourcename]
		[,CONTROL={E|S}]
		[,LOCKOPT={1|2}]
		[,KEEP={NO|YES}]
		[,OWNER={TASK|PARTITION}]

The DLT (Define The Lock) macro generates a control block which is used by the LOCK/UNLOCK macros to enqueue/dequeue a resource access request. The control block, commonly called 'DLT', is generated at the time of program assembly.

NAME=resourcename: specifies the name by which the resource is known to the system for the purpose of access share control. It is by this name, that DOS/VSE controls shared access of the resource as requested by active tasks via the LOCK macro.

These tasks may all be active in one partition, or they may be distributed over several partitions; the resource-share control extends across partitions.

The name may be up to twelve bytes long.

CONTROL={E|S}: defines how the named resource can be shared while your program owns it, which is determined by this specification and your specification for the operand LOCKOPT. A specification of E means the resource is enqueued for exclusive use; a specification of S means the resource is enqueued as sharable.

LOCKOPT={1|2}: This operand, together with the CONTROL parameter, determines how the system controls shared access in response to a LOCK request.
- LOCKOPT=1 and CONTROL=E: no other task is allowed to use the resource concurrently.
- LOCKOPT=1 and CONTROL=S: other 'S' users are allowed concurrent access, but no concurrent 'E' user is allowed.
- LOCKOPT=2 and CONTROL=E: no other 'E' user gets concurrent access; however, other 'S' users can have access to the resource.
- LOCKOPT=2 and CONTROL=S: other 'S' users can have concurrent access and, in addition, one 'E' user is allowed.

KEEP={NO|YES}: This operand may be used to lock the named resource beyond job step boundaries. Only a main task should use this operand. KEEP=NO indicates that the named resource once locked, is to be released automatically at the end of the particular job step. With KEEP=YES, a named resource that is locked remains locked across job steps; it will be automatically released at end-of-job.

OWNER={TASK|PARTITION}: defines whether the named resource, once locked, can be unlocked only by the task which issued the corresponding LOCK request (OWNER=TASK), or whether it can be unlocked by any task within the partition (OWNER=PARTITION).

When OWNER is defined as PARTITION, a LOCK request for the resource must not specify FAIL=WAIT or FAIL=WAITC because deadlock prevention (return code 16) is not supported with OWNER=PARTITION.

DUMP Macro

Name	Operation	Operand
[name]	DUMP	

This macro provides a hexadecimal dump of the following:

- The contents of the entire supervisor area and the used part of the system GETVIS area, or of some supervisor control blocks only (see Note below).
- The contents of the partition that issued the macro.
- The contents of the registers.

Note: The dump includes the contents of some supervisor control blocks only, rather than the entire supervisor area if, the STDOPT job control command specifies DUMP=PART or NO, or if a job control statement //OPTION PARTDUMP or NO-DUMP is submitted.

In addition, the macro causes the job step to be terminated if DUMP was issued by the main (or only) task of the program. If DUMP was issued by a subtask, the macro causes that subtask to be detached without terminating the main task in the partition.

The dump provided by the macro is always directed to SYSLST, which must be opened if disk or tape; if SYSLST is a tape, that tape must be positioned as desired.

If DUMP is issued by a job running in real mode, the storage contents of the partition are dumped only up to the limit as determined by the SIZE parameter of the EXEC job control statement, plus the storage obtained dynamically through the GETVIS macro. If SIZE was not specified, the entire partition will be dumped. If DUMP is issued by a program running in virtual mode, the entire partition is dumped.

ENQ Macro

Name	Operation	Operand
[name]	ENQ	{rcbname|(0)}

A task protects a resource by issuing an ENQ (enqueue) macro. When the RCB, (identified by the rcbname) is enqueued, the task requesting the resource is either queued and executed, or if the requested resource is held by another task, is placed in a wait condition. When the task holding that resource completes, that task issues the DEQ (dequeue) macro. All other tasks that are then waiting for the dequeued resource are freed from their wait condition, and the highest priority task either obtains or maintains control.

If a task is terminated without dequeuing its queued resources, any task subsequently trying to enqueue that resource is abnormally terminated. If a task issues two ENQs without an intervening DEQ for the same resource, the task is canceled. Also, any task that does not control a resource but attempts to dequeue that resource is terminated, unless DEQ appears in the abnormal termination routine. If DEQ

appears in the abnormal termination routine, it is ignored.

Although the main task does not require the program to set up an intertask communication ECB to enqueue and dequeue, every subtask using that facility must have the ECB operand in the ATTACH macro, and that ECB must not be used for any other purpose. Also, a resource can be protected only within the partition containing the ECB.

EOJ Macro

Name	Operation	Operand
[name]	EOJ	

Issue the EOJ macro in the main task or in the only program within a partition, to inform the system that the job step is finished. If a subtask issues an EOJ, the subtask is detached and the remainder of the partition continues. If the main task issues EOJ, all abnormal termination exits (via STXIT AB) are taken for the subtasks still attached.

EXIT Macro

Name	Operation	Operand
[name]	EXIT	{PC|IT|OC|AB|TT|MR}

The EXIT macro is used to return control from your exit control or MR routine to the instruction in your interrupted program immediately after the instruction where the interruption occurred. For AB, control is returned to the instruction following the EXIT AB macro. Your routine is specified in the STXIT macro for MR, in the DTFMR macro. The operands have the following meanings:

PC   Exit from your program check routine.

IT   Exit from your interval timer routine.

OC   Exit from your routine which handles the operator attention interrupt.

AB   Exit from your abnormal task termination routine of your main task.

TT   Exit from your task timer routine.

MR   Exit from your stacker selection routine (MICR document processing) to the the supervisor.

The EXIT macro should be issued only in the corresponding (PC, IT, OC, AB, TT, MR) routine; a program check may occur if this rule is not observed.

Detailed information on the save area and the interrupt status is given in *DOS/VSE Serviceability Aids and Debugging Procedures,*.

For PC, IT, OC, and TT, the interrupt status information and registers are restored from the save area; thus, the save area contents are not over-written.

For AB, the cancel condition and ABEND indication of the affected task are reset. The EXIT AB macro may be used only in main tasks. In a subtask, it would result in an illegal SVC. You have to make sure that the abnormal termination condition has been cleared up by your abnormal task termination routine before using the EXIT AB macro.

FCEPGOUT Macro

You can code the macro in either of the following two formats:

Name	Operation	Operand
[name]	FCEPGOUT	beginaddr,endaddr
		[beginaddr,endaddr]...
[name]	FCEPGOUT	{listname|(1)}

The FCEPGOUT macro causes a specific area in real storage to be paged-out at the next page fault. This request is ignored if the specified area does not contain a full page. This can happen up to an area size of 4K minus 2 bytes (see Figure 5-2).
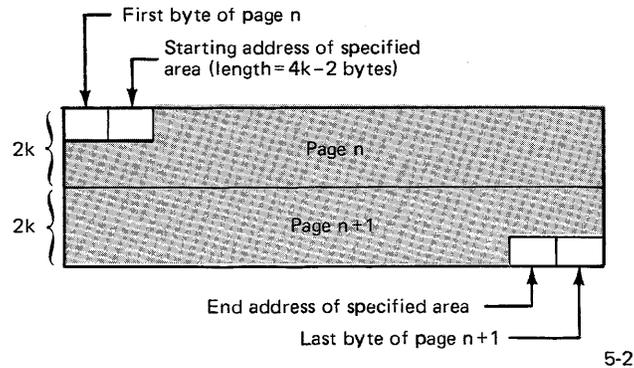
Figure 5-2. Worst case of an area not containing one full page.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation; the return code is set to zero.

beginaddr: Points to the first location of the area to be paged out.

endaddr: Points to the last location of the area to be paged out.

listname|(1): Is the symbolic name of a list of consecutive 8-byte entries as shown below.

X'00'	address constant	length minus 1
0	1	4 7

where:

address constant = Address of the first byte of the area to be paged out.

length = A binary constant indicating the length of the area to be paged out.

A non-zero byte following an entry indicates the end of the list. Register notation may also be used.

Exceptional Conditions

The program is running in real mode.

The page(s) referenced by the macro is (are) outside of the requesting partition.

A page handling request is pending for the referenced page(s).

The page(s) is (are) not in real storage.

The page(s) is (are) fixed.

For those pages the FCEPGOUT request will be ignored.

The supervisor was not generated with PAGEIN=n in the SUPVR macro (in this case the program is canceled).

Return Codes in Register 15

0 All specified pages have been forced for page-out or the request has been ignored because the issuing program is running in real mode.

2 The begin address is greater than the end address, or a negative length has been found.

4 At least one of the requested pages does not belong to the partition in which the issuing program is running. The FCEPGOUT request has only been executed for those pages which belong to the partition of the issuing program.

8 a. At least one of the requested pages is temporarily fixed (via CCW-translation) and/or PFIXed. The FCEPGOUT request has only been executed for the unfixed pages.

b. A page handling request (page fault, temporary fix, PFIX) for at least one of the requested page is pending (caused by asynchronous processing within a partition). The FCEPGOUT request has not been executed for those pages which are involved in a page handling request.

16 List of areas that are to be paged out is not completely in the requesting program's partition. The request is ignored.

Any combination of return codes 0, 2, 4, and 8 is possible.

FETCH Macro

Name	Operation	Operand
[name]	FETCH	{phasename|(S,address)| (1)} [,entrypoint|(S,entrypoint)| (0)] [,LIST= {listname|(S,listname)| (r)}] [,SYS=YES] [,DE=YES] [,MFG= {area|(S,area)|(r2)}]

The FETCH macro loads and gives control to the phase specified in the first parameter from the core image library if this phase is not in the SVA. If the phase is in the SVA, it is not loaded into the partition, but control is given to the phase. For information on how to load phases into the SVA and how to write SVA-eligible (reenterable) phases see *DOS/VSE System Management Guide.*

phasename|(S,address)|(1): For phasename specify the name of the required phase.

If DE=YES is not specified, the address as specified in (S, address) or as loaded into a register points to an 8-byte field that contains the phase name. If DE=YES, the operand has a different meaning; refer to the discussion of the DE operand.

entrypoint|(S,entrypoint)|(0): Control is passed to the address specified by the entrypoint parameter. If this parameter is not specified or invalid, control is passed to the entrypoint determined at link-edit time.

If entrypoint is given in register notation, register 1 must not be used. You preload the register with the entrypoint address.

With S-type notation, the entrypoint is derived from base register and displacement, e.g. (S, offset (reg)). If instead, a symbolic name is used for entrypoint, the macro expansion results in a V-type address constant. The entrypoint does not have to be identified by an EXTRN statement.

LIST= {listname|(S,listname)|(r1)}: For listname specify the name of your local directory list generated in the partition by the GENL macro. When this operand is included, the system scans the local directory list for the required phasename before it initiates a search for this phase name in the pertinent core image library directory.

SYS=YES: If SYS=YES is specified, the system scans the system directory list (SDL) in the SVA and the system core image library before the private core image library (if a private CIL is assigned at all). If

nothing is specified, the private CIL takes precedence.

DE=YES: This operand is useful if your program frequently fetches one specific phase. DE=YES is invalid if LIST is specified.

DE=YES indicates that your program contains a 34-byte field where you have placed a single directory entry like those generated by the GENL macro. If this directory entry is active the directory scan mechanism is bypassed; if not, the entry will be filled in by the supervisor after which it is active.

If the first operand is written as phasename (instead of S-type or register notation) a directory entry will be generated within the macro expansion. The generated directory entry will contain the phasename in the first 8 bytes. If you specify DE=YES and if you use (S, address) or register notation for the first operand, you must set aside the 34-byte field yourself and point to it via this operand. The directory entry must contain the phase name in the first 8 bytes (left-justified and padded with blanks) and X'OB' at displacement X'OB'.

MFG= {area|(S,area)|(r2)}: The operand MFG is required if the program which issues the FETCH macro is to be reenterable. It specifies the address of a 64-byte dynamic storage area, that is, storage which your program obtained through a GETVIS macro. This area is required for system use during execution of the macro.

FREE Macro

Name	Operation	Operand
[name]	FREE	{filename|(1)}

The FREE macro, used in conjunction with the HOLD=YES option of a DTFxx macro, frees a portion of a DASD device that is being held under DASD record (track) protection. On a CKD device, that protected portion is a track; on an FBA device, it is an integral number of contiguous FBA blocks. On an FBA device, the FREE macro is treated as a null operation; all holding and freeing of FBA block ranges is performed implicitly by LIOCS.

The same track (or blocks) can be held more than once without an intervening FREE if the hold requests are from the same task. The same number of FREEs must be issued before the track (or block) is completely freed. However, a task is terminated if more than 16 hold requests are recorded without an intervening FREE, or if a FREE is issued to a file that does not have a hold request for that track (or block). For situations that require the use of the

FREE macro, refer to *DOS/VSE Macro User's Guide*, as listed in the Preface.

{filename|(1)}: This operand is the same as the name specified in the DTF header entry.

FREEVIS Macro

Name	Operation	Operand
[name]	FREEVIS	[ADDRESS={name1|(1)}]
		[,LENGTH={name2|(0)}]
		[,SVA=YES]

The FREEVIS macro releases a block (or blocks) of virtual storage that was obtained by the GETVIS macro.

If you code the macro without any operand, DOS/VSE assumes that the start address of the block to be released is contained in register 1 and that the length of this block was placed into register 0. If the macro is issued without an operand, the macro must not contain a comment unless the comment begins with a comma.

ADDRESS= {name1|(1)}: The start address of the virtual storage block to be released in the GETVIS area may be specified either in a 4-byte field addressed by name1 or in a register.

LENGTH= {name2|(0)}: The length of the virtual storage block to be released may be specified in a 4-byte field addressed by name2 or in a register. The length is specified in number of bytes. The smallest unit of virtual storage that can be released by FREEVIS is (a) 128 bytes if the GETVIS area is part of a partition or (b) 512 bytes if the GETVIS area is part of the SVA. If the specified length is not a multiple of 128 or 512, respectively, it is rounded to the next higher integral multiple of 128 or 512.

SVA=YES: SVA=YES can be specified only in a program that is executed with storage protection key zero. If SVA=YES is specified, DOS/VSE tries to find the block that is to be released in the SVA, otherwise in the pertinent partition.

Return Codes in Register 15

  0 FREEVIS completed successfully
  4 The size of the partition GETVIS area is OK.
  8 The specified length is smaller than zero.
 12 The specified address is not within the GETVIS area or the address is not (a) a multiple of 128 bytes if the GETVIS area is part of a partition, or (b) a multiple of 512 bytes if the GETVIS area is in the SVA.

 16 The specified storage block to be released (ADDRESS+LENGTH) exceeds the GETVIS area.

GENDTL Macro

Name	Operation	Operand
[namd]	GENDTL	[ADDR={name1|(S,name1)|(r1)}]
		[,NAME= {name2|(S,name2)|(r2)}]
		[,CONTROL= {E|S}]
		[,LOCKOPT= {1|2}]
		[,KEEP= {NO|YES}]
		[,OWNER= {TASK|PARTITION}]
		[,LENGTH= {NO|YES}]

The GENDTL macro generates a control block which is used by the LOCK/UNLOCK macros to enqueue/dequeue a resource access request. The control block, commonly called 'DTL' (Define The Lock), is generated at the time of program execution. Space for the DTL is either provided by the program, or it is to be acquired by DOS/VSE.

ADDR= {name1|(S,name1)|(r1)}; specifies either the address or a register pointing to the address where the DTL is to be built. If this operand is omitted, DOS/VSE requests storage to be allocated in the partition's GETVIS area by an implicit GETVIS request. After a successful GETVIS, DOS/VSE places the DTL's address in register 1. Register 15 contains the return code set by the implicit GETVIS request.

NAME= {name2|(S,name2)|(r2)}; specifies the address or a register pointing to the address where an up to 12 byte long resource name is stored. It is by this name that DOS/VSE controls shared access of the resource as requested by active tasks via the LOCK macro. These tasks may all be active in one partition, or they may be distributed over several partitions; the resource-share control extends across partitions.

LENGTH= {NO|YES}; If LENGTH=YES is specified, the GENDTL macro returns the length of the DTL in register 0. With LENGTH=NO, register 0 remains unchanged.

GENL Macro

Name	Operation	Operand
[name]	GENL	phasename1,phasename2, ...
		[{,ADDRESS={area|(S,area)|(r1)}
		,LENGTH=number}]
		[{,ADDRESS= {DYN|DYNAMIC
		[,ERREXIT= {addr|(S,addr)|(r2)}]}]

The GENL macro generates a local directory in the partition. It saves access time if you load the same phases more than once in the course of program execution.

phasename1,phasename2,...: Specify, for these parameters, the names of phases, for which entries in a local directory list are to be built. The list will be generated in alphameric sequence. You may specify up to 200 phase names.

ADDRESS=specification,LENGTH=number: If the ADDRESS operand is omitted, the assembler builds a 34-byte entry within the macro expansion for each of the specified phases and inserts the pertinent phase name in the entry. The rest of the entry is filled in by the supervisor when the phase is called by a FETCH or LOAD macro, with the LIST option for the first time. When, subsequently, the phase is called again, the entry is active.

Coding ADDRESS in conjunction with LENGTH indicates that the directory is to be built, during execution, at a location within your program whose address is given by the ADDRESS operand.

LENGTH gives the length of the field provided for the generation of the directory. If LENGTH is too short the assembler issues an MNOTE.

ADDRESS=DYN,ERREXIT=specification: Coding ADDRESS=DYNAMIC (a short form, ADDRESS=DYN, is allowed) directs the system to acquire, through a GETVIS, as much dynamic storage as needed. Note that in this case the contents of registers 0, 1, 14, and 15 will be over-written by execution of the macro.

ERREXIT is the address of a routine to be executed should the implicit GETVIS fail. If the ERREXIT operand is omitted, an unsuccessful GETVIS will cause the task to be canceled.

GETIME Macro

Name	Operation	Operand
[name]	GETIME	[STANDARD|BINARY|TU] [,LOCAL|,GMT] [,MFG={area|(S,area)|(r)}] .

The GETIME macro obtains the time-of-day at any time during program execution, provided the time of day option was specified at system generation. (If the time of day option was not specified at system generation, issuing GETIME only obtains zeros instead of a valid time.)

STANDARD and LOCAL are assumed if no operands are given. If the macro is issued without an operand, the macro must not contain a comment unless the comment begins with a comma.

As long as no DATE job control statement is supplied, the calendar date and the system date in the communication region are updated every time GETIME is issued. Those dates are therefore accurate at any given moment. However, when the job stream contains a DATE job control statement, only the system date in the communication region is updated when GETIME is used; the calendar date is not changed in that case.

If STANDARD is specified, the time-of-day is returned in register 1 as a packed decimal number of the form *hhmmss*, where *hh* is hours, *mm* is minutes, and *ss* is seconds, with the sign in the low-order half-byte. The time-of-day may be stored and unpacked or edited.

If BINARY is specified, the time-of-day is returned in register 1 as a binary integer in seconds.

If TU is specified, the time-of-day is returned in register 1 as a binary integer in units of 1/300 seconds.

LOCAL\|GMT: Specify LOCAL to obtain the local time or GMT if, in your program, you want to use Greenwich Mean Time.

MFG= {area\|(S,area)\|(r)}: The MFG operand is required if the program is to be reenterable and if option STANDARD applies (with options BINARY or TU, reentrancy is preserved in any case). It specifies the address of a 64-byte dynamic storage area, that is: storage which your program obtained through a GETVIS macro. This area is required for system use during execution of the macro.

GETVIS Macro

Name	Operation	Operand
[name]	GETVIS	[ADDRESS={name1|(1)}] [,LENGTH={name2|(0)}] [,PAGE=YES] [,POOL=YES] [,SVA=YES]

The GETVIS macro retrieves a block of virtual storage from the GETVIS area of your partition or of the SVA. If you code the macro without any operand, DOS/VSE assumes that the length of the desired virtual storage area is contained in register 0 and returns the start address of the area it retrieved for you in register 1. If the macro is issued without an ope-

rand, the macro must not contain a comment unless the comment begins with a comma.

ADDRESS= {name1|(1)}: The start address of the requested virtual storage area is returned by the system either in the 4-byte field specified as a symbol by name1 or in the specified register. (Register 15 must not be used because it contains the return code.) The returned address is valid only if the return code in register 15 is zero. If the operand is omitted, the address is returned in register 1 only.

LENGTH= {name2|(0)}: The length of the requested storage block may be specified either in the 4-byte field (specified as a symbol by name2) or in the specified register. The length is specified in number of bytes. The smallest unit that can be requested by GETVIS is (a) 128 bytes if the GETVIS area is part of a partition or (b) 512 bytes if the GETVIS area is part of the SVA. If the specified length is not a multiple of 128 or 512, respectively, it is rounded to the next higher multiple of 128 or 512. If the operand is omitted, the system assumes that you have specified the length in register 0.

PAGE=YES: If you want the requested storage area to start on a page boundary, specify PAGE=YES. This may reduce the number of page faults.

POOL=YES: If POOL=YES is specified, GETVIS starts searching for the requested virtual storage area at the address specified in register 1. In this case, it is your responsibility to provide a valid address in register 1.

SVA=YES: SVA=YES can be specified only in a program that is executed with storage protection key zero. If SVA=YES is specified, DOS/VSE retrieves the desired block of virtual storage from the SVA. Otherwise, DOS/VSE retrieves the block from the pertinent partition.

Return Codes in Register 15

   0 GETVIS completed successfully.
   4 The size of the partition GETVIS area is OK.
   8 The specified length is negative.
  12 No more virtual storage is available in the GETVIS area.
  32 A hardware (storage) failure occured in the requested real partition GETVIS area.

JDUMP Macro

Name	Operation	Operand
[name]	JDUMP	

This macro provides a hexadecimal dump of the following:

- The contents of either the entire supervisor area and the used part of the system GETVIS area, or of some supervisor control blocks only (see Note below).
- The contents of the partition that issued the macro.
- The contents of the registers.

Note: The dump includes the contents of some supervisor control blocks only, rather than the entire supervisor area, if the STDOPT job control command specifies DUMP=PART or NO, or if a job control statement //OPTION PARTDUMP or NO-DUMP is submitted.

In addition, the macro causes the job to be terminated if JDUMP was issued by the main (or only) task of the program. If JDUMP was issued by a subtask, the macro causes that subtask to be detached without terminating the program in the partition.

The dump provided by the macro is always directed to SYSLST; SYSLST, if disk or tape, must be opened; if SYSLST is a tape, that tape must be positioned as desired.

If JDUMP is issued by a job running in real mode, the storage contents of the partition are dumped only up to the limit as determined by the SIZE parameter of the EXEC job control statement, plus the storage obtained dynamically through the GETVIS macro. If SIZE was not specified, the entire partition will be dumped.

If JDUMP is issued by a program running in virtual mode, the entire partition is dumped.

JOBCOM Macro

Name	Operation	Operand
[name]	JOBCOM	FUNCT= {PUTCOM|GETCOM}, AREA= {address|(r1)}, LENGTH= {length|(r2)}

The JOBCOM macro allows for communication between jobs or job steps in a partition. Information being communicated is stored in a 256-byte area. DOS/VSE provides such an area for each partition. Through the JOBCOM macro, a program either moves information into that area or retrieves information that had previously been stored there by another program. The area remains unaltered from one job (or job step) to the next. Unless it is modified through execution of the JOBCOM macro, the contents of the area remain unchanged over any number of jobs.

The program that issues the JOBCOM macro must provide a register save area 18 fullwords long. Prior

to execution of the macro, register 13 has to point to that save area.

FUNCT={PUTCOM|GETCOM} This operand describes the function that the macro is to perform. Specifying PUTCOM causes information to be stored into the system-supplies area. The number of bytes to be moved is given by the LENGTH operand. If LENGTH yields a value smaller than 256, the remainder of the area is left unaltered.

Specifying GETCOM indicates that information is to be retrieved from the system-supplied area. Again, the number of bytes to be moved is given by the LENGTH operand.

AREA={address|(r1)} This operand gives the address of a field where the program provides (FUNCT=PUTCOM specified) or receives (FUNCT=GETCOM specified) the information to be moved.

LENGTH={length|(r2)} This operand specifies the number of bytes to be moved. The value is either given as a self-defining term or in register notation. If register notation is used, the specified register is expected to contain the length value.

Length should be a positive number up to 256. If it is zero or negative, no information gets moved. If it is greater than 256, only 256 bytes are moved.

LFCB Macro

Name	Operation	Operand
[name]	LFCB	SYSxxx,phasename [,NULMSG] [,FORMS=xxxx][,LPI=n]

The macro can be used to load the forms control buffer (FCB) of a printer dynamically. That printer must not be an IBM 3800 Printing Subsystem; the macro is ignored on an IBM 3800. An FCB whose contents have been changed by means of this macro retains the changed contents until one of the following occurs:

- another LFCB macro is issued for the printer;
- an LFCB command is issued for the printer;
- the SYSBUFLD program is executed to reload the printer's FCB;
- IPL is performed for the system.

The macro, when executed, generates messages to request operator action (such as changing forms), whenever manual action is required, and to inform the operator that the FCB of the specified printer has been reloaded.

Note: If SYSLOG is assigned to a printer, the results of an FCB load operation initiated by an LFCB macro are unpredictable.

SYSxxx: The name of the logical unit associated with the printer whose FCB is to be loaded.

You can specify one of the following:

- SYSLST,
- SYSLOG, or
- a programmer logical unit (SYSnnn) assigned to a printer owned by the partition in which the program is being executed.

phasename: The name by which the phase containing the applicable FCB image is cataloged in the core image library. For information on the contents and format of an FCB image, see the section "System Buffer Load (SYSBUFLD)" in *DOS/VSE System Control Statements.*

NULMSG: This operand specifies that the 80-character verification message, which is normally printed following the FCB load operation, is to be suppressed. This operand, if given, must be specified immediately after phasename.

If this operand is specified, the system continues normal processing immediately after the FCB load operation has been completed, and the operator cannot verify that the proper forms are placed on the printer.

If the operand is omitted, the system prints the last 80 characters of the phase identified by phasename, and causes an additional skip to channel 1.

FORMS=xxxx: This operand specifies the type of forms to be used on the printer whose FCB is being reloaded. For xxxx, a string of up to four alphameric characters can be specified. The specified form number is included in a message to the operator.

LPI=n: This operand, which should not be given for a PRT1-printer (with a device type code of PRT1), specifies the desired number of lines per inch. For n, you can specify either 6 or 8.

If the macro is issued for a PRT1-printer and the specified spacing disagrees with the spacing code in the new buffer image, the system does not execute the FCB load operation and sets the appropriate return code in register 15.

If the macro is issued for a non-PRT1-printer, the system includes the operand in a message to the operator.

Return Codes in Register 15

Successful completion of the FCB load operation is indicated to the problem program by a return code of 0. Note, however, that for an IBM 3800, register 15 contains 0, although the macro was not executed. If the operation fails, register 15 contains one of the return codes listed below; in this case the FCB retains its original contents. The return codes are:

Return Code Dec	Hex	Meaning
4	X'04'	The assigned printer is a PRT1-printer, and the LPI operand specified in the macro disagrees with the FCB image.
8	X'08'	No LUB is available for the specified logical unit.
12	X'0C'	The specified logical unit has not been assigned or is currently unassigned.
16	X'10'	The specified logical unit is assigned to a device without an FCB.
20	X'14'	The printer assigned to the specified logical unit is down.
24	X'18'	The specified FCB image phase has not been found.
28	X'1C'	The specified FCB image phase is invalid for the printer assigned to the specified logical unit.

By testing register 15, you can determine in your program whether or not the operation has failed. If the operation has failed, you can either terminate the job step or continue processing. Should you decide to continue processing, then the system bypasses the execution of the LFCB macro.

LOAD Macro

Name	Operation	Operand
[name]	LOAD	{phasename|(S,address)|(1)} [,loadpoint|(S,loadpoint)|(0)] [,LIST= {listname|(S,listname)|(r1)}] [,SYS=YES] [,DE=YES] [,TXT=NO] [,MFG= {area|(S,area)|(r2)}]

The LOAD macro loads the phase specified in the first parameter from the core image library (if this phase is not in the SVA) and returns control to the calling phase. After execution of the macro, the entry-point address of the called phase is returned to you in register 1. For a non-relocatable phase, this address is the entry-point determined at link-edit time. For a relocatable phase, the entry point is adjusted by the relocation factor. If the phase is in the SVA, it is not loaded; the entry point address in the SVA, however, is returned in register 1.

phasename|(S,address)|(1): For phasename specify the name of the required phase.

If DE=YES is not specified, the address as specified in (S, address) or as loaded into a register points to an 8-byte field that contains the phase name. If DE=YES, the operand has a different meaning; refer to the discussion of the DE operand.

loadpoint|(S,loadpoint)|(0): If loadpoint is provided the load-point address specified to the linkage editor is overridden, and the phase is loaded at the specified address. The address used must be outside the supervisor area. When an overriding address is given, the entry-point address is relocated and returned in register 1. An overriding loadpoint address must not be specified for a phase that had been linked as a member of an overlay structure.

If the phase is non-relocatable, none of the other addresses in the phase are relocated; if the phase is relocatable, however, the entry point and address constants are updated with the relocation factor.

If loadpoint is given in register notation, the register used must not be register 1. Preload the register with the loadpoint address.

With (S,...) notation, the loadpoint address is derived from base register and displacement as assembled for loadpoint in the (S,loadpoint) specification.

LIST= {listname|(S,listname)|(r1)}: For listname specify the name of your local directory list generated in the partition by the GENL macro. When this operand is included, the system scans the local directory list for the name of the required phase before it initiates a search for this phase name in the pertinent core image library directory.

SYS=YES: If SYS=YES is specified, the system scans the system directory list (SDL) in the SVA and the system core image library before the private core image library (if a private CIL is assigned at all). If nothing is specified, the private CIL takes precedence.

DE=YES: This operand is useful if your program frequently loads one specific phase. DE=YES is invalid if LIST is specified.

DE=YES indicates that your program contains a 34-byte field where you have placed a single directory entry like those generated by the GENL macro. If this directory entry is active the directory scan mechanism is bypassed; if not, the entry will be filled in by the supervisor and then becomes active.

If the first operand is written as phasename (instead of S-type or register notation) a directory entry will be generated within the macro expansion.

The generated directory entry will contain the phasename in the first 8 bytes.

If you use (S,address) or register notation for the first operand you must set aside the 34-byte field yourself and point to it via this operand. The directory entry must contain the phase name in the first 8 bytes (left-justified and padded with blanks) and X'0B' at displacement X'0B'.

TXT=NO: The specification TXT=NO (with LIST=listname or DE=YES) is useful if a phase is loaded more than once in the course of your program. It causes a search for the directory entry without transfer of the contents (or text) of the phase itself and it indicates in the directory entry if and where the phase was found. This can be used to accomplish either of the following:

1. The directory entry can be filled in from the core image library for later FETCH/LOAD calls without the overhead of text transfer.

2. You can establish whether a given phase is present in a core image library or the SVA since register 0 contains the address of the directory entry and byte 16 of the directory entry appears as:

X'06'	if the phase is not found
X'12'	if the phase is in the SVA
X'0A'	if the phase is in the private CIL.

Note: Test for these conditions by means of a Test Under Mask (TM) instruction, not a Compare instruction.

MFG= {area|(S,area)|(r2)}: The operand MFG is required if the program which issues the LOAD macro is to be reenterable. It specifies the address of a 64-byte dynamic storage area, that is, storage which your program obtained through a GETVIS macro. This area is required for system use during execution of the macro.

LOCK Macro

Name	Operation	Operand
[name]	LOCK	{name|(S,name)|(r)}
		[,FAIL={RETURN|WAITC|WAIT}]

The LOCK macro enqueues the task for accessing the named resource. The resource must have been defined in the program by a DTL (Define The Lock) control block. A DTL is generated by issuing a DTL or GENDTL macro; it may be modified by issuing a MODDTL macro.

{name|(S,name)|(r)}: specifies the DTL address.

FAIL= {RETURN|WAITC|WAIT}: defines the

system action in case the resource cannot be obtained:

- *FAIL=RETURN:* causes the system to return control back to the requesting program in any case. The requesting program has to check the return code in register 15 to fine out whether or not the request was successful.
- *FAIL=WAITC:* causes the system to place the requesting task in the wait state if the requested resource is found to be locked by another task. In all cases, control returns to the requesting program. The requesting program has to check the return code in register 15 to fine out whether or not the request was successful, or whether an error occurred.
- *FAIL=WAIT:* requests the system to return control to the requesting task when the resource can be obtained. If the resource is locked by another task, the requesting task is set into the wait state until the resource is freed. In case of an error condition (return codes 12, 16, 20), the requesting task is canceled.

WAIT or WAITC cannot be specified if the resource is defined with OWNER=PARTITION.

Return Codes in Register 15:

0 Successful request: the resource is locked for the task (or for the partition if the resource is defined with partition ownership).

4 Resource not available: the resource is already locked with a locking status that allows no concurrent access.

8 The lock table space is exhausted.

12 The lock request is inconsistent with previous lock requests (by the same or other tasks).

16 The request would have resulted in a deadlock condition.

20 DTL format error.

Figure 5-3 presents a summary of system actions by return codes, depending on the specification of the FAIL operand.

Figure 5-4 summarizes how DOS/VSE controls access to a resource, depending on the satisfactions

R.C. =	LOCK FAIL =		
	RETURN	WAITC	WAIT
0	RETURN	RETURN	RETURN
4	RETURN	WAIT	WAIT
8	RETURN	RETURN	WAIT
12	RETURN	RETURN	CANCEL
16	RETURN	RETURN	CANCEL
20	RETURN	RETURN	CANCEL

Figure 5-3. System actions by return code and FAIL operand.

Control definition in requesting DTL		Control definition in owning DTL			
		LOCKOPT=1		LOCKOPT=2	
		CONTROL=S	CONTROL=E	CONTROL=S	CONTROL=E
LOCKOPT=1	CONTROL=S	G	W	I	I
	CONTROL=E	W	W	W	W
LOCKOPT=2	CONTROL=S	I	W	G	G
	CONTROL=E	I	W	G	W

G	The lock request is granted (return code 0).
I	The lock request is inconsistent with current lock status (return code 12).
W	Access to the resource cannot be granted (return code 4 or 16).

Figure 5-4: System action depending on control definition in DTLs.

of the CONTROL and LOCKOPT operands in the DTL or GENDTL macro. The illustration assumes that a task issues a LOCK request for a resource which is already locked.

A task or partition may lock a resource more than once. DOS/VSE maintains a lock request count for the resource.

When a resource is defined with LOCKOPT=1, a task may issue up to 255 LOCK requests with CONTROL=S. When a resource is defined with LOCKOPT=2, up to 255 LOCK requests with CONTROL=S and (if no other task locks the resource, exclusively) one LOCK request with CONTROL=E are allowed.

When a resource is locked more than once by a task, this task has to issue at least as many UNLOCK requests as it issued LOCK requests before it gives up the resource completely. If the resource is defined with OWNER=PARTITION, the unlocking may be done by any task in the partition.

MODDTL Macro

Name	Operation	Operand
[name]	MODDTL	ADDR={name1|(S,name1)|(r1)}
		[,NAME={name2|(S,name2)|(r2)}]
		[,CONTROL={E|S}]
		[,LOCKOPT={1|2}]
		[,KEEP={NO|YES}]
		[,OWNER={TASK|PARTITION}]
		[,CHANGE={ON|OFF}]

The MODDTL macro modifies operands (fields) of a DTL (Define The Lock) control block. A DTL is used by the LOCK/UNLOCK macros to enqueue/dequeue a specific resource. The control block must have been generated by the DTL or GENDTL macro.

Operands not specified in the MODDTL macro leave the corresponding field in the DTL unchanged. There are no default values for the MODDTL macro.

ADDR={name1|(S,name1)|(r1)}: specifies the address of the DTL.

NAME={name2|(S,name2)|(r2)}: specifies the address or a register pointing to the address where an up to 12-byte long resource name is stored. It is by this name, that DOS/VSE controls shared access of the resource as requested by active tasks via the LOCK macro. These tasks may all be active in one partition, or they may be distributed over several partitions; the resource-share control extends across partitions.

CHANGE={ON|OFF}: CHANGE=ON sets up the DTL such that a subsequent UNLOCK macro with name1 as operand would not release the resource, rather keep the resource locked, but reduce its locking status. Reducing the lock status can be done only when the current lock status is defined with strongest possible values: CONTROL=E and LOCKOPT=1. At least one of the operands CONTROL and LOCKOPT should be specified too. CHANGE=OFF causes a subsequent UNLOCK macro to resume its normal function: to dequeue the resource.

For a description of the other parameters refer to the DTL macro.

MVCOM Macro

Name	Operation	Operand
[name]	MVCOM	to,length,{from|(0)}

The MVCOM macro modifies the content of bytes 12 through 23 of the communication region of the partition from which the macro is issued. This area is commonly referred to as the user area.

The following example shows how to move three bytes from the symbolic location DATA into bytes 16 through 18 of the communication region:

MVCOM 16,3,DATA

to: specifies the address (relative to the first byte of the region) of the first communication region byte to be modified.

length: represents the number of bytes (1 to 12) to be inserted.

from|(0): represents the address (either as a symbol or in register notation) of the bytes to be inserted.

PAGEIN Macro

You can code the macro in either of the following two formats:

Name	Operation	Operand
[name]	PAGEIN	beginaddr,endaddr [,beginaddr,endaddr]... [,ECB={ecbname|(0)}]
[name]	PAGEIN	{listname|(1)} [,ECB={ecbname|(0)}]

The PAGEIN macro causes specific areas to be brought into real storage before their contents are needed by the requesting program. If the requested area is already in real storage the attached page frame will get low priority for the next page-outs. This function, however, does not include any fixing, so that it cannot determine whether all areas requested will still be in real storage when the entire request has been completed.

On a system that was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation. If the ECB operand is specified, the ECB will be posted.

beginaddr: Points to the first byte of the area to be paged in.

endaddr: Points to the last byte of the area to be paged in.

listname|(1): Is the name of a list of consecutive 8-byte entries as shown below.

X'00'	address constant	length minus 1	
0	1	4	7

where:

address constant = Address of the first byte of the area to be paged in.

length = A binary constant indicating the length of the area to be paged in.

A non-zero byte following an entry indicates the end of the list.

ECB=ecbname|(0): Specifies the name of the ECB, a fullword defined by your program, which is to be posted when the operation is complete. An invalid ECB address causes the task to be canceled.

Return Information

The return information can be obtained from the ECB, byte 2. The meaning of these bits is shown below.

Bit	Meaning if bit is one:
0	PAGEIN request is finished.
1	The page table is full, the request cannot be queued at this time for further handling; the request is ignored, bit 0 is set.
2	One or more of the requested pages are outside the requesting program's partition; PAGEIN is not performed for these pages.
3	At least one negative length has been detected in the area specifications; PAGEIN is not performed for these areas.
4	List of areas that are to be paged in is not completely in the requesting program's partition; the request is ignored, bit 0 is set.
5	Paging activity is too high in the system, no performance improvement is possible.
6-7	Reserved

Any combination of the return bits in the ECB is possible.

Use the WAIT macro with the ecbname as operand for completion of the PAGEIN macro, before the bits in byte 2 of the ECB are tested.

PDUMP Macro

Name	Operation	Operand
[name]	PDUMP	{address1|(r1)}, {address2|(r2)} [,MFG={area|(S,area)|(r3)}]

This macro provides a hexadecimal dump of the general registers and of the virtual storage area contained between the two address expressions (address1 and address2). The contents of registers 0 and 1 are over-written, but the CPU status is retained. Thus, PDUMP furnishes a dynamic dump (snapshot) useful for program checkout. Processing continues with your next instruction.

The dump is always directed to SYSLST with 121-byte records. The first byte is an ASA control character. When SYSLST is a disk drive, you must issue an OPEN or OPENR macro to any DTF assigned to SYSLST after each PDUMP that is executed. The OPEN or OPENR macro updates the disk address maintained in the DTF table to agree with the address where the PDUMP output ends. If the OPEN or OPENR is not issued, the address is not updated, and the program is canceled when the next PUT is issued.

If non-addressable areas were included in the range of PDUMP, a message will be printed to indicate this.

{address1|(r1)}, {address2|(r2)}: One or both of the addresses can be specified in register notation. If address2 is not greater than address1, or address1 is greater than the highest address in the allocated virtual storage, the macro results in no operation. If the value in address2 is greater than the end of the allocated virtual storage area, the virtual storage between address1 and the end of the allocated virtual storage is dumped.

MFG= {area|(S,area)|(r3)}: The MFG operand is required if the program which issues the PDUMP is to be reenterable. It specifies the address of a 64-byte dynamic storage area, that is, storage which you obtained by GETVIS macro; this area is needed by the system during execution of the macro.

PFIX Macro

You can code the macro in either of the following two formats:

Name	Operation	Operand
[name]	PFIX	beginaddr,endaddr [,beginaddr,endaddr]...
[name]	PFIX	{listname|(1)}

The PFIX macro causes specific pages to be brought into real storage and fixed in their page frames until they are released at some later time. The maximum number of pages that may be fixed at any one time is specified via the ALLOCR command. Each time a page is fixed a counter for that page is incremented. This counter may never exceed 255 for any page.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation; the return code is set to zero.

beginaddr: Points to the first byte of the area to be fixed.

endaddr: Points to the last byte of the area to be fixed.

listname|(1): Is the name of a list of consecutive 8-byte entries as shown below.

X'00'	address constant	length minus 1
0	1	4 7

where:

address constant   =   Address of the first byte of the area to be fixed.

length   =   A binary constant indicating the length of the area to be fixed.

A non-zero byte following an entry indicates the end of the list. Register notation may be used.

Exceptional Conditions

If a PFIX causes the count of fixes for a page to exceed 255, the task issuing the PFIX is canceled.

If it is not possible to fix all pages requested, then none will be fixed.

If PFIX is issued in a program running in real mode, it is ignored and register 15 contains 0.

Return codes in Register 15

0 if the pages were successfully fixed.
4 if the number of pages to be fixed for one request exceeds the number of PFIXable page frames; in order for this PFIX request to be satisfied, more PFIXable storage must be allocated through the ALLOCR command.
8 if not enough page frames are available in the partition because of previous PFIXes or current system resource usage; this PFIX request could, however, be satisfied at another time without reallocating PFIXable storage.
12 if one of the specified addresses was invalid.

PFREE Macro

You can code the macro in either of the following two formats:

Name	Operation	Operand
[name]	PFREE	beginaddr,endaddr [,beginaddr,endaddr]...
[name]	PFREE	{listname|(1)}

Each page in the virtual address area is assigned a 'PFIX counter'. If a page is not fixed - that is, if it is subject to normal page management - the counter is 0. Whenever a page is fixed by using a PFIX macro its counter is increased by one. All pages whose counters are greater than 0 remain fixed in real storage.

The PFREE macro decrements the counter of a specified page by 1. If a PFREE is issued for a page whose counter is 0, that PFREE is ignored since the page has already been freed.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation; the return code is set to zero.

beginaddr: Points to the first byte of the area to be freed.

endaddr: Points to the last byte of the area to be freed.

listname|(1): Is the symbolic name of a list of consecutive 8-byte entries as shown below.

X'00'	address constant	length minus 1	
0	1	4	7

where:

address constant	=	Address of the first byte of the area to be fixed.
length	=	A binary constant indicating the length of the area to be fixed.

A non-zero byte following an entry indicates the end of the list.

Exceptional Conditions
If PFREE is issued by a program running in real mode, the macro is ignored.

Return codes in Register 15
0 if the pages were successfully freed.
12 if one of the addresses specified was invalid.

POST Macro

Name	Operation	Operand
[name]	POST	{ecbname|(1)} [,SAVE= {savearea|(0)}]

This macro provides intertask communication by posting an ECB (it turns on byte 2, bit 0). A POST

issued to an ECB removes a task waiting for the ECB from the wait state.

ecbname|(1): Provides the address of the ECB that is to be posted.

SAVE= {savearea|(0)}: This operand may be used for taking a specific waiting subtask out of the wait state. The operand causes DOS/VSE to locate the save area whose address is specified in the operand and to take only the subtask associated with this save area out of the wait state. This task normally is waiting for the specified ECB to be posted.

Although time is saved by specifying this operand, other tasks waiting for this ECB are not taken out of the wait state for this event by this issuance of the POST macro. This does not guarantee that they will stay in the wait state until another POST is issued. On the contrary, other events could cause the other tasks to be dispatched. For this reason the POST macro should not be used with the SAVE operand to control subtask operation unless separate ECB's are used. Otherwise, it should be used only to save time. When a POST is issued without the SAVE operand, all tasks waiting for the ECB are taken out of the wait state, and the highest priority task regains control.

RCB Macro

Name	Operation	Operand
[name]	RCB	

The RCB macro generates an 8-byte word-aligned resource control block (RCB); this block allows you to protect a user-defined resource if the ENQ macro is issued before (and the DEQ macro is issued after) each use of the resource. The format of the RCB and its use is shown below.

Bytes	Purpose of bits
0	All bits are set to 1 to indicate that the resource has been placed in a priority queue by the ENQ macro.
1-3	Reserved
4	If bit 0=1: Another task is waiting to use the resource. Bits 1-7: Reserved.
5-7	ECB address of current resource owner.

REALAD Macro

Name	Operation	Operand
[name]	REALAD	{address|(1)}

In S/370 mode, the REALAD macro returns the real address corresponding to a specified virtual address. If issued in ECPS:VSE mode, the macro returns the specified virtual address.

address|(1): Is the virtual address to be converted. It can be specified as a symbol or in register notation.

Register 0 returns the address corresponding to the specified virtual address if and only if the virtual address points to a PFIXed page, otherwise register 0 contains 0. Thus, the macro can be used to test if a page is PFIXed.

Note: The pages of a partition running in real mode are treated as if they were fixed.

RELEASE Macro

Name	Operation	Operand
[name]	RELEASE	(SYSnnn[,SYSnnn]...) [,savearea]

RELEASE specifies the names of programmer logical units to be released. RELEASE may be used only for units used within a given partition .

The operand SYSnnn specifies the programmer logical unit that is to be released. Up to 16 units may be specified in a list, which must be enclosed in parentheses.

All the units specified are checked by the assembler to assure that no system logical units are requested for release. If system logical units are specified, an MNOTE is issued and such units are ignored. Before any release is attempted, a check is made for ownership of the unit. If the requesting partition does not own the unit, or if the unit is already unassigned, the request is ignored.

savearea: Is the name of an 8-byte word-aligned area where registers 0 and 1 are saved for your program. If the operand is not provided, the contents of registers 0 and 1 are over-written.

The macro expansion includes a unit table and loads the table's address into register 0. If the savearea operand is specified, the macro expansion saves registers 0 and 1.

If there is no permanent assignment, the device is unassigned. If the device is at permanent assignment level, no action is taken on the unit.

Recommendation: You should inform the system operator via a message that the assignment was released.

RELPAG Macro

You can code the macro in either of the following two formats:

Name	Operation	Operand
[name]	RELPAG	beginaddr,endaddr [,beginaddr,endaddr]...
[name]	RELPAG	{listname|(1)}

The RELPAG macro causes the contents of one or more storage areas to be released. If the affected areas are in real storage when the RELPAG macro is executed, their contents are not saved but are over-written when the associated page frames are needed to satisfy pending page frame requests.

After the RELPAG macro has been executed for an area and a location in that area is referenced again during the current program execution, the related page is attached to a page frame which contains all zeros.

The storage area is released only if it contains at least one full page. You can be sure of this only if the specified area is 4K minus 1 byte, or bigger (see Figure 5-5).
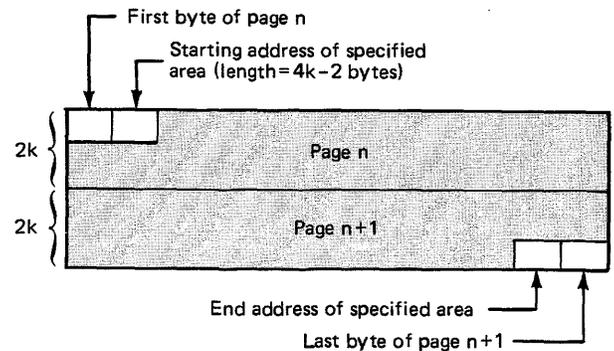

Figure 5-5. Worst case of an area not containing one full page.

beginaddr: Points to the first byte of the area to be released.

endaddr: Points to the last byte of the area to be released.

listname|(1): Is the symbolic name of a list of consecutive 8-byte entries as shown below.

X'00'	address constant	length minus 1
0	1	4	7

where:

address constant   = Address of the first byte of the
                     area to be released.
length   = A binary constant indicating
                     the length of the area to be re-
                     leased.

A non-zero byte following an entry indicates the end
of the list.

Register notation may be used.

Exceptional Conditions
* The program is running in real mode.
* The area is, fully or partially, outside of the
  virtual partition of the requesting program.
* A page handling request is pending for the ref-
  erenced page(s).
* The page(s) is (are) fixed. For these pages, the
  RELPAG request will be ignored.

* The supervisor was not generated with
  PAGEIN=n in the SUPVR macro (in this case the
  program will be canceled).

Return Codes in Register 15
0 All referenced pages have been released or the
   request has been ignored because the request-
   ing program is running in real mode.
2 The begin address is greater than the end ad-
   dress, or a negative length has been found.
4 The area, fully or partially, does not belong to
   the partition where the issuing program is run-
   ning. The release request has only been execut-
   ed for those pages which belong to the partition
   of the issuing program.
8 a. At least one of the requested pages is tempo-
   rarily fixed (via CCW-translation) and/or
   PFIXed. The release request has only been exe-
   cuted for the unfixed pages.
   b. A page handling request (page fault, tempo-
   rary fix, PFIX) for at least one of the requested
   pages is pending (caused by asynchronous
   processing within a partition). The release re-
   quest has not been executed for those pages
   which are involved in a page handling request.
16 List of areas that are to be released is not com-
   pletely in the requesting program's partition.
   The request is ignored.

Any combination of the return codes 2, 4, and 8 is
possible.

RETURN Macro

Name	Operation	Operand
[name]	RETURN	(r1[,r2])

The RETURN macro restores the registers whose
contents were saved and returns control to the call-
ing program.

The operands r1,r2 specify the range of the regis-
ters to be reloaded from the save area of the pro-
gram that receives control. The operands are written
as self-defining values. When inserted in an LM
machine instruction, the operands cause the desired
registers in the range from 14 through 12 (14, 15, 0
through 12) to be restored from words 4 through 18
of the save area. If r2 is omitted, only the register
specified by r1 is restored. To access this save area,
register 13 must contain the save area address.
Therefore, the address of the save area should be
loaded into register 13 before execution of the
RETURN macro.

RUNMODE Macro

Name	Operation	Operand
[name]	RUNMODE	

The RUNMODE macro returns the following inform-
ation to the program issuing it:

* Register 1 contains 0 if the issuing program is
  running in virtual mode.
* Register 1 contains 4 if the issuing program is
  running in real mode.

No operand is required for this macro.

SAVE Macro

Name	Operation	Operand
[name]	SAVE	(r1[,r2])

The SAVE macro stores the contents of specified
registers in the save area provided by the calling
program.

The operands r1,r2 specify the range of the regis-
ters to be stored in the save area of the calling pro-
gram. The address of this area is passed to the pro-
gram in register 13. The operands are written as
self-defining values so that they cause desired regis-
ters in the range of 14 through 12 (14, 15, 0 through
12) to be stored when inserted in an STM assembler
instruction.

Registers 14 and 15, if specified, are saved in
words 4 and 5 of the save area. Registers 0 through
12 are saved in words 6 through 18 of the save area.
The contents of a given register are always stored in
a particular word in the save area. For example,
register 3 is always saved in word 9 even if register 2
is not saved.

If r2 is omitted, only the register specified by r1 is saved.

SETIME Macro

Name	Operation	Operand
[name]	SETIME	{timervalue|(1)},[tecbname|(r)][,PREC]

The SETIME macro sets the interval timer to the specified value.

timervalue|(1): The amount of time for the interval. This value can be specified either as an absolute expression or in register notation. If register notation is used, the pertinent register must contain the time value.

The largest allowable value is 55,924 seconds — equivalent to 15 hours, 32 minutes, 4 seconds — if PREC is omitted, and 8,388,607 — equivalent to 7 hours, 46 minutes, 2 seconds (approximately) — if PREC is specified. The positional operand PREC indicates that the timer value is expressed in units of 1/300 of a second. When PREC is omitted, the timer value is in seconds.

tecbname|(r): Specifies the name (address if register notation is used) of a timer event control block (TECB) which must have been defined previously in your program by a TECB macro. If you use register notation, register 0 and 1 must not be used. After having executed the SETIME macro, the system returns the TECB address in register 1.

When a program is restarted from a checkpoint, any timer interval set by a SETIME macro is not restarted.

The SETIME macro must not be used within an abnormal termination exit routine.

The setting of a time interval can be utilized in one of two ways. In the first method, the program sets up linkage to an interval timer exit routine by issuing the STXIT IT macro. When the time interval specified in SETIME elapses, control is given to that routine (if a routine is not supplied to the supervisor by the time of the interruption, the interruption is ignored). When using SETIME in this way, the tecbname operand must not be specified.

In the second method, specifying the tecbname operand causes a TECB to be posted when the time interval has elapsed: bit 0 of byte 2 in the TECB is set to 1. This event bit is set to 0 when SETIME is issued. A task may be waiting for the TECB to be posted by having issued the WAIT or WAITM macro.

SETPFA Macro

Name	Operation	Operand
[name]	SETPFA	{entryaddress|(0)}

The SETPFA macro either establishes or terminates linkage to a page fault appendage routine that is to be entered each time a page fault occurs or is completed. You will find more information on how to write such a routine in *Appendix F* of the *DOS/VSE Macro User's Guide*, as listed in the Preface.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation.

entryaddress|(0): If an entry address is specified, execution of the macro establishes linkage to the appendage routine. The routine at that address will be entered every time a page fault in the associated task occurs or is satisfied. The routine to be entered and all areas referenced by the routine must be fixed in real storage using the PFIX macro before SETPFA is issued. The entry address may be specified as a symbol or in register notation.

If SETPFA is issued without an operand, the linkage to the page fault appendage is terminated. Each issuance of SETPFA supersedes all previous SETPFA's for that task.

A page fault appendage is called only when a page fault occurs in the task owning the appendage. If a page fault occurs while a supervisor service routine is working for the owning task, the appendage is not called. The same may apply to an IBM-supplied component such as ACF/VTAME.

SETT Macro

Name	Operation	Operand
[name]	SETT	{timervalue|(1)}

The SETT macro sets the task timer to the value, in milliseconds, specified in the operand. The largest allowable value is 21474836 milliseconds. A register can be specified, and if it is, that register must contain the number of milliseconds in binary. You can use the SETT macro only if your supervisor was generated with TTIME=partition-ID specified in the FOPT generation macro.

The SETT macro can be issued only by the main task of the partition owning the task timer. If it is issued by a program running in a partition not owning the task timer, the program is canceled and an error message indicating illegal SVC is printed.

SETT must not be used within an abnormal termi-
nation exit routine.

The time interval is decremented only while the
task is executing. When the specified time interval
has elapsed, the task timer routine supplied in the
STXIT TT macro is entered.

If a routine is not supplied to the supervisor by
the time of the interruption, the interrupt is ignored.
When a program is restarted from a checkpoint,
timer intervals set by a SETT macro are not restarted.

STXIT Macro

Name	Operation	Operand
To establish linkage:		
[name]	STXIT	{AB|IT|PC|OC|TT}
		,{rtnaddr|(0)} , {savearea|(1)}
		[,OPTION= {DUMP|NODUMP}]
To terminate linkage:		
[name]	STXIT	{AB|IT|PC|OC|TT}

The STXIT (set exit) macro establishes or terminates
linkage from the supervisor to an exit routine of
your program for handling the specified condition.
Linkage must be established before an interrupt
occurs. Use the EXIT macro to return from these
routines.

When restarting a program from a checkpoint,
any STXIT linkages established prior to the check-
point are destroyed.

If, in an exit routine, you are issuing I/O
request(s) requiring the same logic module as your
main routine, you must generate a read-only module
by specifying RDONLY=YES in the DTF and in the
logic module. Both the main routine and the exit
routine require a save area of their own. Detailed
information on the save area and interrupt status is
given in *DOS/VSE Serviceability Aids and Debug-
ging Procedures.*

AB: An abnormal task termination routine is en-
tered if a job or task is terminated for some reason
other than a CANCEL, DETACH, DUMP, or EOJ macro
being issued by the task itself. Upon entry to the
task's abnormal termination routine,

- termination messages and a partition dump are
  produced, depending on selected options (see
  the OPTION operand below).
- bit 1 of byte 2 in the task's attachment ECB is
  posted (if an ECB was specified in the ATTACH
  macro).
- register 0 contains the abnormal termination
  code in its low order byte (see Figure 5-6).

- register 1 points to the task's abnormal-
  termination save area, which contains the inter-
  rupt status information and the contents of reg-
  isters 0 through 15 at the time of abnormal ter-
  mination.

The abnormal termination routine can then ex-
amine this data and take whatever action is neces-
sary.

Macros which might be used in this routine are,
for instance, POST and CLOSE. However, if an ab-
normal termination condition occurs *within* an ab-
normal termination routine, the job or task is abnor-
mally terminated without regard to an abnormal
termination exit. Thus, your program's abnormal
termination routine should avoid macros such as
ENQ, CHKPT, and any I/O macros which may cause
an abnormal termination.

After the appropriate action is taken, your abnor-
mal termination routine should end with a CANCEL,
DETACH, DUMP, or EOJ macro.

If your routine issues the DUMP macro, the system
produces a storage map of the partition even if job
control option NODUMP was specified. For the parti-
tion, SYSLST may be assigned to a 3211 printer. If, in
addition, indexing was used before your abnormal
termination routine received control, a certain num-
ber of characters on every line of the printed dump
may be lost, unless you reload the printer's FCB
(forms control buffer) by issuing an LFCB macro
before you issue the DUMP macro. The FCB image
to be loaded in this case must not have an indexing
byte.

If the CANCEL macro is issued in the main task,
the entire partition is terminated with every subtask
abnormal termination exit taken in order of priority.

If the system was generated with the multitasking
option, each task may require its own abnormal
termination routine. A main task can attach a sub-
task with an ABSAVE operand. This assumes the
subtask will use the main task's abnormal termina-
tion routine. However, the subtask may override
this specification by issuing its own STXIT AB macro.

If an abnormal termination condition occurs and
linkage has not been established to an abnormal
termination routine, processing in the partition is
abnormally terminated. However, if the abnormal
termination condition occurs in a subtask without
exit linkage, only the subtask is terminated.

When subtasks are detached or canceled, associ-
ated time intervals and exit linkages are cleared.

IT: An interval timer interruption routine is entered
when the specified interval elapses. If the program

Hexadecimal representation	Specific abnormal termination code meaning
00	Default value for all cases other than those listed below
0F	Invalid FBA DASD address for SYSFIL
10	Normal EOJ
11	No channel program translation for unsupported device
12	Insufficient buffer space for channel program translation
13	CCW with count greater than 32K
14	Page pool too small
15	Page fault in disabled program (not a supervisor routine)
16	Page fault in MICR stacker select or page fault appendage routine
17	Main task issued a CANCEL macro with subtask still attached
18	Main task issued a DUMP macro with subtask still attached
19	Operator replied cancel as the result of an I/O error message
1A	An I/O error has occurred (see interrupt status information)
1B	Channel failure
1C	CANCEL ALL macro issued in another task
1D	Main task terminated with subtask still attached
1E	A DEQ macro was issued for a resource but tasks previously requesting a resource cannot be found because their save areas (containing register 0) were modified
1F	CPU failure
20	A program check occurred
21	An invalid SVC was issued by the problem program or macro
22	Phase not found in the core image library
23	CANCEL macro issued
24	Canceled due to an operator request
25	Invalid virtual storage address given (outside partition)
26	SYSxxx not assigned (unassigned LUB code)
27	Undefined logical unit
28	Reserved
29	Reserved
2A	I/O error on page data set
2B	I/O error during fetch from private core image library
2C	Page fault appendage routine passed illegal parameter to supervisor
2D	Program cannot be executed/restarted due to a failing storage block
2E	Invalid resource request (possible deadlock)
2F	More than 255 PFIX requests for one page
30	Read past a /& statement
31	I/O error queue overflow during system error recovery procedure
32	Invalid DASD address
33	No long seek on a DASD
35	Job control open failure
36	Page fault in I/O appendage routine
38	Wrong privately translated CCW
39	Reserved
40	ACF/VTAME error, invalid condition
41	ACF/VTAME error, invalid condition
42	Invalid extent information violates DASD file protection
FF	Unrecognized cancel code

Figure 5-6. Abnormal termination codes.

issuing the STXIT macro instruction is a ACF/VTAME application program, the interruption exit will not be taken while ACF/VTAME is processing any request on behalf of the application program. The exit will be taken when ACF/VTAME has completed the program's request.

An interval timer interruption is ignored if no exit linkage has been established.

If an interval timer interrupt occurs while an interval timer exit routine is still processing, the handling of the interrupt is delayed. When processing ends with EXIT IT, the IT exit routine is entered again to process the new IT interrupt. (This can only occur if a short time interval was issued in your exit routine).

OC: An operator communication interruption routine is entered when you press the request key on the console and type the MSG command. In case of multitasking, only the main task can process this condition.

An operator communication interruption is ignored if no exit linkage has been established.

PC: A program check interruption routine is entered when a program check occurs. If a program check occurs in a routine being executed from the logical transient area, the job containing the routine is abnormally terminated.

A program check interruption routine can be shared by more than one task within a partition. To accomplish this, issue the STXIT macro in each subtask with the same routine address but with separate save areas. To successfully share the same PC routine, the routine must be reenterable, that is, it must be capable of being used concurrently by two or more tasks. (The specified exit is not taken if the program check occurs while ACF/VTAME is processing a request issued by the program.)

If a program check condition occurs in a main task without exit linkage, processing in the partition is terminated. However, if this same condition occurs in a subtask, only the subtask is terminated.

TT: Linkage to a task timer interruption routine can be established only if TTIME=partition-ID was specified in the FOPT macro for supervisor assembly.

A task timer interruption routine is entered when the time interval specified in the SETT macro has elapsed. The STXIT (and EXIT) TT macro can be issued only by the main task of the partition owning the task timer. If it is issued by a program running in a partition not owning the task timer, the program

is canceled and an error message indicating illegal SVC is printed.

A task timer interruption is ignored if no exit linkage has been established. A task timer interrupt is ignored if it occurs while a task timer exit routine is still processing. (This can happen only if a short time interval was issued in your exit routine).

rtnaddr: Entry point address of the routine that processes the condition described in the first operand. Your exit routine may be located anywhere in the program.

savearea: Address of a 72-byte area in which the supervisor stores the old interrupt status information and general registers 0 through 15, in that order. Your program must have a separate save area for each routine that is included.

OPTION= {DUMP|NODUMP}: This operand can be used only when setting up linkage (STXIT AB) to an abnormal termination exit routine. It determines whether termination messages and a dump will be issued upon entry to the routine.

If the OPTION operand is omitted or if OPTION=DUMP is specified, termination messages are issued upon enry to the abnormal termination routine. In addition, a partition dump is produced unless the job control option NODUMP is active.

If OPTION=NODUMP is specified, neither a termination message nor a dump is produced. However, if the abnormal termination routine terminates abnormally, termination messages and the dump are given regardless of the OPTION specification in the STXIT macro.

If your routine ends with a DUMP macro and the STXIT macro was specified without OPTION=NODUMP, you will obtain two dumps.

Figure 5-7 shows what happens when one of the five conditions occurs while an STXIT routine is being processed within a particular partition.

TECB Macro

Name	Operation	Operand
[name]	TECB	

The TECB macro generates a timer event control block which can be referred to by the symbol you specify in the name field. This block contains an event bit that indicates when the time interval specified in SETIME has elapsed. The format of this block is as follows:

Byte	Purpose of bits	
0-1	Reserved	
2	0: If 0 =	time specified in SETIME has not elapsed.
	If 1 =	time specified in SETIME has elapsed.
	1-7: Reserved	
3	Reserved	

TESTT Macro

Name	Operation	Operand
[name]	TESTT	[CANCEL]

The TESTT macro can be used only if TTIME=partition ID was specified in the FOPT generation macro for supervisor assembly.

The TESTT macro is used to test the amount of time that has elapsed from a task timer interval set by an associated SETT. The TESTT macro returns the time remaining in the interval, expressed in hundredths of milliseconds in binary, in register 0.

If CANCEL is specified, the remaining time of the interval is canceled, and the task timer exit routine is

Routine being Processed	Condition Occurring				
	AB	IT	OC	PC	TT
AB	C	D	I	C	D
IT	S	E	H	H	H
OC	S	H	I	H	H
PC	S	H	H	T	H
TT	S	H	H	H	I

C   Job canceled immediately without entering AB routine again.

D   Interrupt is delayed and the TT or IT exit routine is entered after the EXIT AB macro is issued. If no EXIT AB is issued, the interrupt is ignored.

E   Handling of new timer interrupt delayed until execution of EXIT IT for original interrupt.

H   Condition honored. When processing of new routine completes, control returns to interrupted routine.

I   Condition ignored.

S   Execution of the routine being processed is suspended, and control transfers to the AB routine.

T   Job abnormally terminated. If AB routine is present, its exit is taken. Otherwise, a system abnormal termination occurs.

Notes:

1. If an operator communication interruption routine or a program check interruption routine is in process when a timer interrupt occurs, your timer routine will be processed; when it completes, control returns to interrupted routine.

2. If a task is using a logical transient routine when a timer interrupt occurs, your timer routine is not entered until the logical transient routine is released.

Figure 5-7. Effect of an AB, IT, OC, PC, or TT interrupt while an STXIT routine is being executed.

not entered. If the macro is issued without an operand, the macro must not contain a comment unless the comment begins with a comma.

The TESTT macro can be issued only by the main task of the partition owning the task timer. If it is issued by a program running in a partition not owning the task timer, the program is canceled and an error message indicating illegal SVC is printed.

TPIN Macro

Name	Operation	Operand
[name]	TPIN	

The TPIN macro is available primarily for the telecommunication applications that require immediate system response. The macro causes one or more partitions (other than the one issuing the macro) to be deactivated. The number of partitions that can be deactivated is specified in the TPBAL command. The partitions to be deactivated are the ones with the lowest priorities. This request is ignored in each of the following cases:

- The operator has not made TP balancing active by means of the TPBAL command.
- None of the partitions specified in the TPBAL command contains a program running in virtual mode.
- The only partition that could be affected by TP balancing is the partition that issued the TPIN request.
- There is no paging in the system.

The TPIN macro must always be used in conjunction with the TPOUT macro. The operand field is ignored.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation.

TPOUT Macro

Name	Operation	Operand
[name]	TPOUT	

The TPOUT macro causes DOS/VSE to reactivate partitions that had been deactivated by the TPIN macro.

Failure to issue the TPOUT macro can cause considerable and unnecessary performance degradation in the batch partition(s). The operand field is ignored.

If your supervisor was generated with VM=YES (in the SUPVR generation macro), execution of the macro results in a null operation.

TTIMER Macro

Name	Operation	Operand
[name]	TTIMER	[CANCEL]

The TTIMER macro is used to test how much time has elapsed of an interval which was set in the same task by the associated SETIME macro. The TTIMER macro returns the time remaining of the interval, expressed in hundredths of seconds in binary, in register 0.

If CANCEL is specified, the time interval set in that task is canceled. As a result of the TTIMER CANCEL macro, the interval timer interruption routine of the task (to which linkage may have been extablished by an STXIT IT macro) does not receive control. If the associated SETIME macro specified the same name of a TECB, that TECB's event bit is set on. If the macro is issued without an operand, the macro must not contain a comment unless the comment begins with a comma.

UNLOCK Macro

Name	Operation	Operand
[name]	UNLOCK	{{name|(S,name)|(r)}|ALL}

The UNLOCK macro dequeues the issuing task (or partition) from the named resource, to which the task had previously been enqueued via the LOCK macro. The resource must have been defined to the system by a DTL or GENDTL macro.

In addition, the UNLOCK macro can be used to lower the lock control level. Reduction of the lock control level may be done only if the issuing task is currently locked, onto the resource, with the most stringent control level possible: CONTROL=E and LOCKOPT=1 (the CONTROL and LOCKOPT parameters are described with the DTL macro). The resource then continues to be held by the task; however, another task waiting for this resource can be dispatched again and may or may not gain shared access (see also the description of the LOCK macro). In order to use the UNLOCK macro for this purpose, the MODDTL macro must have been issued with the CHANGE operand set ON.

{name|(S,name)|(r)}: specifies the DTL address.

ALL: frees all resources which are locked by the task and whose DTLs were defined with KEEP=NO. If UNLOCK ALL is issued by the main task, not only the resources locked by that task are unlocked, but also those which have been locked by subtasks, with OWNER=PARTITION specified for DTL generation.

Return Codes in Register 15:

0 Successful request; the resource has been unlocked.
4 The resource is not locked for the unlocking task.
8 DTL format error.

The UNLOCK ALL macro does not provide a return code; register 15 remains unchanged.

VIRTAD Macro

Name	Operation	Operand
[name]	VIRTAD	{address|(1)}

In S/370 mode, the VIRTAD macro returns the virtual address corresponding to a specified real address.

address|(1): Is the real storage address to be converted. It can be given as a symbol or in register notation.

In ECPS:VSE mode, only a virtual address can be specified as input. The macro returns that same address. Register 0 returns the virtual address only if:

* for S/370 mode, the specified real address points to a page frame that contains a PFIXed page,
* for ECPS:VSE mode, the specified virtual address points to a PFIXed page.

Otherwise register 0 contains 0. Thus, the macro can be used to test if a page is PFIXed.

Note: The pages of a program running in real mode are considered to be fixed.

WAIT Macro

Name	Operation	Operand
[name]	WAIT	{ecbname|(1)}

With the WAIT macro a task sets itself into the wait state until the specified event control block (ECB) is posted (bit 0 of byte 2 turned on). Each task of the system can wait on any ECB within the same partition. Do not use the WAIT macro for waiting on a telecommunication ECB, an RCB, or a CCB other than one that is associated with an I/O operation started in the same task.

When a WAIT macro is processed in a multiprogramming or multitasking environment, control is given to the supervisor, which makes the CPU available to another task in the same or another partition.

One example of an ECB to be posted is the timer event control block (TECB) specified in the SETIME macro. The task issuing the WAIT remains in the wait state until the timer interval has elapsed (event bit in the TECB turned on).

WAITM Macro

Name	Operation	Operand
[name]	WAITM	{ecb1,ecb2,...|listname|(1)}

The WAITM macro enables your program or task to wait for one of a number of events to occur. Control returns to the task when at least one of the ECBs specified in the WAITM macro is posted.

The operand provides the addresses of the ECBs to be waited upon. The names of ecb1, ecb2... are assumed when at least two operands are supplied. A maximum of 16 names can be coded. If one operand is supplied, it is assumed to be the name (listname) of a list of consecutive full-word addresses that point to the ECBs to be waited upon. The first byte following the last address in the list must be nonzero to indicate the end of the list.

When control returns to a waiting task, register 1 points to the posted ECB (byte 2, bit 0 set on). Any fullword can be used as ECB if bit 0 of byte 2 of that word indicates a completed event. Examples of these blocks are CCBs and TECBs. However, a task never regains control if it is waiting for a CCB to be posted by another task's I/O completion. A MICR CCB gets posted only when the device stops, not when a record is read. Furthermore, telecommunication ECBs and all RCBs must not be waited for because their format does not satisfy a WAIT or a WAITM (that is, bit 0 of byte 2 would not be posted).

XECBTAB Macro

Name	Operation	Operand
[name]	XECBTAB	TYPE={DEFINE|DELETE| CHECK|RESET|DELETALL} ,XECB=xecbname [,XECBADR={xecbfield| (S,xecbfield)|(r1)}] [,ACCESS={XPOST|XWAIT}] [,MFG={area|(S,area)|(r2)}]

You can use the XECBTAB macro only if XECB=YES was specified in the FOPT generation macro for supervisor assembly. The XECBTAB macro can be used (1) to define, for the specified cross-partition event control block (XECB), an entry in the supervisor's XECB table, (2) to delete such an entry, (3) to check for the presence of an entry, or (4) to reset an entry.

An XECB for which an entry has been defined to the DOS/VSE supervisor can be referred to by XPOST and XWAIT macros; an XECB can be referred to also by a WAIT or WAITM macro if the task issuing the macro has previously defined the XECB (with ACCESS=XWAIT).

TYPE= {DEFINE|DELETE|CHECK| RESET|DELETALL}

The operand specifies the type of operation to be performed:

DEFINE causes a new XECB table entry to be defined to the supervisor.

DELETE causes an entry to be deleted from the supervisor's XECB table. TYPE=DELETE can be specified only for an XECB for which an entry has been defined previously in the same program.

CHECK causes DOS/VSE to check whether or not an entry for a specific XECB has been defined already. If that entry is present, specifying CHECK causes the address of both the XECB and the associated XECB table entry to be returned.

RESET causes DOS/VSE to clear the information in the supervisor XECB table that indicates which task communicates with the program having defined the XECB. TYPE=RESET can be specified only for an XECB for which an entry has been previously defined in the same program.

After RESET, any task can attempt to establish a new connection with the owner. With ACCESS=XPOST, however, if the task currently connected to the XECB is issuing an XWAIT macro (at the time of RESET), this task will probably establish connection again, nullifying the RESET operation.

DELETALL performs three actions:

- Causes all entries in the XECB table that were defined previously by the issuing task to be deleted.

- Breaks the communication between any XECB owner and the issuing task (that is, clears the information in the XECB table that indicates that the issuing task communicates with the XECB owner).

- Post all tasks as ready-to-run that are waiting for an XPOST by the issuing task. If these tasks are XWAITing on the XECB, they will get a return code of X'08'.

Note: If TYPE=DELETALL is specified, the XECB and ACCESS operands must not be specified. XECBTAB with TYPE=DELETALL specified does not provide a return code; all registers remain unchanged.

XECB=xecbname: specifies the name of the XECB. If the XECBADR operand is not present, xecbname is

the symbolic address of the 4-byte (or larger) XECB field. If, however, XECBADR is specified, xecbname is the name by which the control block is known between partitions; the symbolic address of the control block field is given by XECBADR.

The XECB field must be defined in your program, except when TYPE=CHECK is specified: in that case, the XECB field may be defined in another program.

XECBADR= {xecbfield|(S,xecbfield)|(r1)}: XECBADR is used only if TYPE=DEFINE is specified. It provides the symbolic address of the 4-byte (or larger) field that is to be used as XECB.

ACCESS= {XPOST|XWAIT}: This operand can be used together with TYPE=DEFINE to specify whether the program will be allowed to post the XECB or wait for another program to do the posting.

XPOST is assumed if the operand is omitted. It specifies that the program will be allowed to post the XECB . Specifying XPOST implies that only one other active task is allowed to issue an XWAIT macro against the XECB.

XWAIT specifies that the program will be allowed to wait for one other task to post the XECB.

MFG= {area|(S,area)|(r2)}: The MFG operand is required if the program that issues the XECBTAB macro is to be reenterable. The operand specifies the address of a 64-byte storage area, that is, storage which your program obtains by a GETVIS macro. This area is needed for use by the system during execution of the macro.

The MFG operand is only useful in conjunction with XECBADDR coded in either of the two notations (S,xecbfield) or (r1).

Feedback Information

Figure 5-8 shows the return codes that are supplied by DOS/VSE in register 15. The illustration also indicates whether or not DOS/VSE returns the addresses of the pertinent XECB and the associated table entry in registers 1 and 14, respectively.

	X'00'	X'04'	X'08'
DEFINE	XECB named is stored in the table *	XECB named is already in the table **	The XECB table is full **
DELETE	XECB named is removed from the table **	XECB named was not found **	The issuing program did not define the XECB **
CHECK	XECB named was found in the table *	XECB named was not found **	N/A
RESET	XECB named communication bytes cleared **	XECB named was not found **	The issuing program did not define the XECB **

* Register 1 contains the address of the XECB and register 14 contains the address of the table entry.

** Registers 1 and 14 are set to zero.

Figure 5-8. XECBTAB feedback information.

XPOST Macro

Name	Operation	Operand
[name]	XPOST	XECB= {xecbname|(1)}, POINTRG=(14)

You can use the XPOST macro only if XECB=YES was specified in the FOPT generation macro for supervisor assembly.

The XPOST macro provides for cross-partition communication by posting the specified XECB (the macro sets bit 0 of byte 2 to 1). An XPOST macro issued against an XECB causes the task waiting for this XECB to be removed from the wait state (the waiting task may have issued an XWAIT, WAIT or WAITM with a previously defined XECB). This task may have been activated in the same or in another partition.

If the XPOST macro is used in a main-line loop, the macro should be preceded by a test which ensures that the other partition's task waiting for the event that is being posted must receive control and execute the function for which this event is a prerequisite.

To perform this test, a second XECB needs to be defined. This XECB allows the originally waiting task in its main-line loop to post completion of its function as an event for which the originally posting task must wait.

Resetting bit 0 of byte 2 of the XECB is a user responsibility.

Once a task has issued an XPOST macro for an
XECB (with ACCESS=XWAIT), no other task can issue
an XPOST for this XECB, until the connection is end-
ed.

XECB={xecbname|(1)}: specifies the name of the
XECB to be posted. The name you specify must be
the same as the one used to define the XECB. If reg-
ister notation is used, the specified register must
point to an 8-byte character field that contains the
XECB name left-justified and padded with blanks.
Do not specify 14 or 15 if you choose to use ordinary
register notation.

POINTRG=(14): specifies the register that points
to the XECB table entry associated with the named
XECB. Do not specify register 1 or 15 if you choose
to use ordinary register notation.

To obtain the address of the associated XECB ta-
ble entry, issue earlier in the program an XECBTAB
macro for the same XECB and with TYPE=CHECK or
TYPE=DEFINE specified. When DOS/VSE executes
the XECBTAB macro, the system returns, in register
14, the address of the pertinent XECB table entry.
Figure 5-9, which shows a coding example for the
use of the XWAIT macro, applies to the XPOST macro
accordingly.

Note that if the POINTRG register contains 0 (or
any invalid value), all entries in the XECB table are
searched to determine the correct address; no error is
indicated.

Return Codes
When DOS/VSE returns control to the issuing task,
register 15 contains one of the following return
codes:

Return Code	Meaning
X'00'	Successful completion. The named XECB has been posted.
X'04'	The named XECB has no associated table entry in the XECB table.
X'0D'	The XECB referred to in the XPOST macro was defined with ACCESS=XPOST specified in the XECBTAB macro, but the task that issued the XPOST macro does not own this XECB.
X'0E'	The XECB referred to in the XPOST macro was defined with ACCESS=XWAIT specified in the XECBTAB macro and either (1) the task that issued the XPOST macro also defined the XECB or (2) the XECB has been posted previously during the same execution by another task.

Note: Following the execution of an XPOST macro, register 1
and 14 are set to zero.

XWAIT Macro

Name	Operation	Operand
[name]	XWAIT	XECB={xecbname|(1)}, POINTRG=(14)

You can use the XWAIT macro only if XECB=YES was
specified in the FOPT macro for supervisor assembly.

The XWAIT macro enables the issuing task to wait
for an XECB to be posted by another task that is exe-
cuting in the same or in another partition. Control
returns to the issuing task when the XECB is posted
or if DOS/VSE detects an error condition.

Once a task has issued an XWAIT macro for an
XECB (with ACCESS=XPOST) to be posted, no other
task can issue an XWAIT for this XECB, until the con-
nection is ended.

XECB={xecbname|(1)}: specifies the name of the
XECB, which may be defined in the same or another
program. The name you specify must be the same as
the one used to define the XECB. If register notation
is used, the specified register must point to an 8-byte
field that contains the name of the XECB left-
justified and padded with blanks. Do not specify
register 14 or 15 if you choose to use ordinary regis-
ter notation.

POINTRG=(14): specifies the register that points
to the XECB table entry associated with the named
XECB. Do not specify register 1 or 15 if you choose
to use ordinary register notation.

To obtain the address of the associated XECB ta-
ble entry, issue earlier in the program an XECBTAB
macro for the same XECB and with the TYPE=CHECK
or TYPE=DEFINE specified. When DOS/VSE executes
the XECBTAB macro, the system returns, in register
14, the address of the pertinent XECB table entry.
Figure 5-9 shows a coding example for the use of the
XWAIT macro; in that example, the required contin-
uation character is not shown. The example assumes
that the XECB was defined by a program executing
in another partition by (source) instructions as fol-
lows:

```
           .
           .
           XECBTAB    TYPE=DEFINE,
                      XECB=MYECB
           .
           .
           .
   MYECB   DC         F'0'
           .
           .
```

```
                    .
                    .
                    .
WA1TLP    XECBTAB    TYPE=CHECK,
                     XECB=MYECB
          LTR        15,15
          BNZ        ERROR
          LA         1,XECBNAME
          XWAIT      XECB=(1),
                     POINTRG=(14)
                    .
                    .
                    .
XECBNAME  DC         CL8'MYECB
```

Figure 5-9. Coding example showing the use of XECBTAB with TYPE=CHECK and of XWAIT.

Note that if the POINTRG register contains 0 (or any invalid value), all XECBs are searched to determine the correct address; no error is indicated.

Return Codes

When DOS/VSE returns control to the issuing task, register 15 contains one of the following return codes:

Return code	Meaning
X'00'	Successful completion. The named XECB has been posted.
X'04'	The named XECB has no associated table entry in the XECB table.
X'08'	The other task using this XECB has broken communication without issuing an XPOST.
X'0D'	The XECB referred to in the XWAIT macro was defined with ACCESS=XWAIT specified in the XECBTAB macro, but the task that issued the XWAIT macro does not own this XECB.
X'0E'	The XECB referred to in the XWAIT macro was defined with ACCESS=XPOST specified in the XECBTAB macro and either (1) the task that issued the XWAIT macro also defined the XECB or (2) the task did not define the XECB, but another task is already waiting for the XECB to be posted.

Note: Following the execution of an XWAIT macro, registers 1 and 14 are set to zero.

DOS/VSE
Macro Reference
GC24-5140-0

READER'S
COMMENT
FORM

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM may use or 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.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.*

	Yes	No
• Does the publication meet your needs?	☐	☐
• Did you find the material:		
Easy to read and understand?	☐	☐
Organized for convenient use?	☐	☐
Complete?	☐	☐
Well illustrated?	☐	☐
Written for your technical level?	☐	☐

• What is your occupation? _____

• How do you use this publication:

As an introduction to the subject?	☐	As an instructor in class?	☐
For advanced knowledge of the subject?	☐	As a student in class?	☐
To learn about operating procedures?	☐	As a reference manual?	☐

Your comments:

If you would like a reply, please supply your name and address on the reverse side of this form.

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments.)

GC24-5140-0

Reader's Comment Form

||| |||
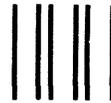
NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS   PERMIT NO. 40   ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department G60
P. O. Box 6
Endicott, New York 13760

Fold   Fold

If you would like a reply, *please print:*

Your Name _____

Company Name _____ Department _____

Street Address _____

City _____

State _____ Zip Code _____

IBM Branch Office serving you _____

IBM ®

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N. Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N. Y., U. S. A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N. Y., U. S. A. 10601

Cut or Fold Along Line

GC24-5140-0

IBM®