




DIGITAL
RESEARCH™

Concurrent CP/M™
Operating System
**Programmer's
Reference Guide**



DIGITAL
RESEARCH™

Concurrent CP/M™
Operating System
Programmer's
Reference Guide



COPYRIGHT

Copyright ©1984 by Digital Research Inc. All rights reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written permission of Digital Research, Post Office Box 579, Pacific Grove, California, 93950.

DISCLAIMER

Digital Research makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Further, Digital Research reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Digital Research to notify any person of such revision or changes.

TRADEMARKS

CP/M and CP/M-86 are registered trademarks of Digital Research. ASM-86, Concurrent CP/M, DDT, DDT-86, MP/M, MP/M-86, and PL/I are trademarks of Digital Research. Intel and MCS are registered trademarks of Intel Corporation. ISIS-II is a trademark of Intel Corporation. IBM is a registered trademark of International Business Machines.

The *Concurrent CP/M Operating System Programmer's Reference Guide* was printed in the United States of America.

First Edition: January 1984

Foreword

Concurrent CP/M[®] is a multi- or single-user operating system targeted specifically for the Intel[®] 8086/8088/80186 family of microprocessors. It supports multiple CP/M programming environments each implemented on a virtual console. A different task runs concurrently in each environment.

This manual describes the invariant programming interface to Concurrent CP/M. It supports the applications programmer who must create applications programs that run in the Concurrent CP/M environment.

Section 1 offers an overview of the entire operating system.

Section 2 describes the structure of the Concurrent CP/M file system.

Section 3 explains the format, structure, and uses of transient commands in the Concurrent CP/M environment.

Section 4 explains the creation of transient command files in the Concurrent CP/M environment.

Section 5 documents the structure and creation of resident system processes or resident command files permanently installed in the Concurrent CP/M environment.

Section 6 describes all the Concurrent CP/M system calls.

Concurrent CP/M is supported and documented through four manuals:

- *The Concurrent CP/M Operating System User's Guide* (hereinafter cited as *Concurrent CP/M User's Guide*) documents the user's interface to Concurrent CP/M, explaining the various features used to execute applications programs and Digital Research utility programs.
- *The Concurrent CP/M Operating System Programmer's Reference Guide* (hereinafter cited as *Concurrent CP/M Programmer's Reference Guide*) documents the applications programmer's interface to Concurrent CP/M, explaining the internal file structure and system entry points, information that is essential for creating applications programs that run in the Concurrent CP/M environment.

- The *Concurrent CP/M Operating System Programmer's Utilities Guide* (hereinafter cited as *Programmer's Utilities Guide*) documents the Digital Research utility programs that programmers use to write, debug, and verify applications programs written for the Concurrent CP/M environment.
- The *Concurrent CP/M Operating System System Guide* (hereinafter cited as *Concurrent CP/M System Guide*) documents the internal, hardware-dependent structures of Concurrent CP/M.

Table of Contents

1 Concurrent CP/M System Overview

1.1	Introduction.....	1-1
1.2	Supervisor (SUP)	1-5
1.3	Real-time Monitor (RTM)	1-5
1.3.1	Process Dispatching	1-5
1.3.2	Queue Management.....	1-7
1.3.3	System Timing Functions	1-8
1.4	Memory Module (MEM)	1-9
1.5	Basic Disk Operating System (BDOS).....	1-9
1.6	Character I/O Module (CIO).....	1-10
1.7	Virtual Console Screen Management	1-10
1.8	Extended Input/Output System (XIOS)	1-11
1.9	Terminal Message Processes (TMP)	1-12
1.10	Transient Programs	1-12
1.11	System Call Calling Conventions	1-12
1.12	SYSTAT: System Status	1-13

2 The Concurrent CP/M File System

2.1	File System Overview	2-1
2.1.1	File-access System Calls.....	2-2
2.1.2	Drive-related System Calls	2-3
2.2	File Naming Conventions.....	2-5
2.3	Disk Drive and File Organization.....	2-8
2.4	File Control Block Definition.....	2-9
2.4.1	FCB Initialization and Usage	2-12
2.4.2	File Attributes	2-14
2.4.3	Interface Attributes	2-16
2.5	User Number Conventions.....	2-17
2.6	Directory Labels and XFCBs	2-18
2.7	File Passwords	2-22
2.8	File Date and Time Stamps: SFCBs	2-24
2.9	File Open Modes	2-26
2.10	File Security	2-27
2.11	Extended File Locking.....	2-30
2.12	Compatibility Attributes	2-31
2.13	Multisector I/O	2-34

Table of Contents (continued)

2.14	Concurrent File Access	2-35
2.15	File Byte Counts	2-37
2.16	Record Blocking and Deblocking	2-38
2.17	Reset, Access, and Free Drive	2-39
2.18	BDOS Error Handling	2-43
3 Transient Commands		
3.1	Transient Program Load and Exit	3-1
3.1.1	Shared Code	3-2
3.1.2	8087 Support	3-2
3.1.3	8087 Exception Handling	3-3
3.2	Command File Format	3-3
3.3	Base Page Initialization	3-5
3.4	Parent/Child Relationships	3-8
3.5	Direct Video Mapping	3-8
4 Command File Generation		
4.1	Transient Execution Models	4-1
4.1.1	The 8080 Memory Model	4-2
4.1.2	The Small Memory Model	4-4
4.1.3	The Compact Memory Model	4-5
4.2	GENCMD	4-6
4.3	Intel Hexadecimal File Format	4-9
5 Resident System Process Generation		
5.1	Introduction to RSPs	5-1
5.2	RSP Memory Models	5-1
5.2.1	8080 Model RSP	5-2
5.2.2	Small Model RSP	5-2
5.3	Multiple Copies of RSPs	5-3
5.3.1	8080 Model	5-3
5.3.2	Small Model	5-4
5.3.3	Small Model with Shared Code	5-4
5.4	Creating and Initializing an RSP	5-4
5.4.1	The RSP Header	5-7

Table of Contents (continued)

5.4.2	The RSP Process Descriptor	5-8
5.4.3	The RSP User Data Area	5-9
5.4.4	The RSP Stack	5-9
5.4.5	The RSP Command Queue	5-9
5.4.6	Multiple Processes within an RSP	5-10
5.5	Developing and Debugging an RSP	5-11
6 System Calls		
6.1	System Call Summary	6-13
6.2	Concurrent CP/M System Calls	6-20
6.2.1	Console I/O System Calls	6-21
6.2.2	Device System Calls	6-41
6.2.3	Disk Drive System Calls	6-44
6.2.4	File-access System Calls	6-64
6.2.5	List Device I/O System Calls	6-122
6.2.6	Memory System Calls	6-128
6.2.7	Process/Program System Calls	6-139
6.2.8	Queue System Calls	6-162
6.2.9	System Information System Calls	6-174

Appendixes

A	System Call Summary by Function Number	A-1
B	ASCII and Hexadecimal Conversions	B-1
C	Error Codes	C-1
D	ECHO.A86 Listing	D-1
E	8087 Exception Handling	E-1
	Glossary	Glossary-1
	Index	Index-1

Table of Contents (continued)

Tables

1-1.	Registers Used by System Calls	1-13
2-1.	File System Calls	2-3
2-2.	Valid Filename Delimiters	2-6
2-3.	Filetype Conventions	2-7
2-4.	Drive Capacity	2-8
2-5.	FCB Field Definitions	2-11
2-6.	File Attribute Definitions	2-15
2-7.	BDOS Interface Attributes F5' and F6'	2-16
2-8.	Directory Label Field Definitions	2-19
2-9.	XFCB Field Definitions	2-21
2-10.	Password Protection Modes	2-22
2-11.	Compatibility Attribute Definitions	2-32
2-12.	BDOS Physical Errors	2-44
2-13.	BDOS Extended Errors	2-45
2-14.	BDOS Error Codes	2-47
2-15.	BDOS Physical and Extended Errors	2-49
3-1.	Group Descriptors	3-4
3-2.	Group Descriptor Fields	3-4
4-1.	Concurrent CP/M Memory Models	4-1
4-2.	Intel Hex Field Definitions	4-11
6-1.	System Call Categories	6-2
6-2.	Concurrent CP/M System Calls	6-4
6-3.	System Call Summary	6-13
6-4.	Data Structures Index	6-18
6-5.	CX Error Code Reports	6-19
6-6.	ACB Field Definitions	6-22
6-7.	C_RAWIO Calling Values	6-31
6-8.	Console Buffer Field Definitions	6-34
6-9.	C_READSTR Line-editing Characters	6-34
6-10.	DPB Field Definitions	6-49
6-11.	PFCB Field Definitions	6-87
6-12.	FCB Initialization	6-89
6-13.	MCB Field Definitions	6-129
6-14.	MPB Field Definitions	6-130

Table of Contents (continued)

Tables

6-15.	APB Field Definitions	6-140
6-16.	Command Line Buffer Field Definitions	6-143
6-17.	PD Field Definitions	6-147
6-18.	UDA Field Definitions	6-152
6-19.	CPB Field Definitions	6-160
6-20.	QPB Field Definitions	6-163
6-21.	QD Field Definitions	6-169
6-22.	SYSDAT Table Data Fields	6-180
6-23.	TOD Field Definitions	6-186
A-1.	System Call Summary by Function Number	A-1
B-1.	ASCII Symbols	B-1
B-2.	ASCII Conversion Table	B-1
C-1.	Concurrent CP/M Error Codes	C-1

Figures

1-1.	Concurrent CP/M Virtual/Physical Environments	1-1
1-2.	Concurrent CP/M Functional Modules	1-3
2-1.	FCB - File Control Block	2-10
2-2.	Directory Label Format	2-18
2-3.	XFCB - Extended File Control Block	2-20
2-4.	Directory Record with SFCB	2-24
2-5.	SFCB Subfields	2-24
2-6.	Disk System Reset	2-41
3-1.	CMD File Header Format	3-3
3-2.	Group Descriptor Format	3-3
3-3.	Concurrent CP/M Base Page Values	3-6
4-1.	Initial Program Stack	4-2
4-2.	Concurrent CP/M 8080 Memory Model	4-3

Table of Contents (continued)

4-3.	Concurrent CP/M Small Memory Model	4-4
4-4.	Concurrent CP/M Compact Memory Model	4-5
4-5.	Intel Hexadecimal File Formats	4-10
5-1.	8080 and Small RSP Models	5-2
5-2.	RSP Head Format	5-3
5-3.	RSP Command Queue Message	5-5
5-4.	RSP Data Segment	5-7
6-1.	ACB - Assign Control Block	6-21
6-2.	Console Buffer Format	6-33
6-3.	Drive, R/O, or Login Vector Structure	6-44
6-4.	DPB - Disk Parameter Block	6-48
6-5.	Disk Free Space Field Format	6-63
6-6.	PFCB - Parse Filename Control Block	6-86
6-7.	MCB - Memory Control Block	6-128
6-8.	MPB - Memory Parameter Block	6-129
6-9.	MFPB - M_FREE Parameter Block	6-132
6-10.	APB - Abort Parameter Block	6-139
6-11.	CLI Command Line Buffer	6-142
6-12.	PD - Process Descriptor	6-146
6-13.	UDA - User Data Area	6-151
6-14.	CPB - Call Parameter Block	6-159
6-15.	QPB - Queue Parameter Block	6-163
6-16.	QD - Queue Descriptor	6-168
6-17.	BDOS Version Number Format	6-174
6-18.	BIOS Descriptor Format	6-175
6-19.	Operating System Version Number Format	6-176
6-20.	SERIAL Number Format	6-177
6-21.	SYSDAT Table	6-179
6-22.	TOD Time-of-Day Structure	6-185

Table of Contents (continued)

Listings

6-1.	Memory Control Block Definition	6-129
6-2.	Memory Parameter Block Definition	6-130
6-3.	Queue Parameter Block Definition.....	6-164
D-1.	ECHO.A86.....	D-1
E-1.	8087 Exception Handling.....	E-2

Section 1

Concurrent CP/M System Overview

1.1 Introduction

Concurrent CP/M is a multi- or single-user, multitasking operating system that lets you run multiple programs simultaneously by initiating tasks on two or more terminals or virtual consoles. Applications programs have access to system calls used by Concurrent CP/M to control the multiprogramming environment. As a result, Concurrent CP/M supports extended features, such as communication among and synchronization of independently running processes. Figure 1-1 depicts the relationships between applications programs, virtual environments, virtual consoles, and the user terminal.




Figure 1-1. Concurrent CP/M Virtual/Physical Environments

In the Concurrent CP/M environment there is an important distinction between a program and a process. A program is simply a block of code residing somewhere in memory or on disk; it is essentially static. A process, on the other hand, is a dynamic entity. You can think of it as a logical machine that executes not only the program code, but also the operating system routines necessary to support the program's functions.

When Concurrent CP/M loads a program, it creates a process associated with the loaded program. Subsequently, it is the process, rather than the program, that obtains access to the system's resources. Thus, Concurrent CP/M monitors the process, not the program. This distinction is a subtle one, but vital to your understanding of system operation as a whole.

Processes running under Concurrent CP/M fall into two categories: transient processes and Resident System Processes (RSPs). Transient processes run programs loaded into memory from disk in response to a user command or system calls made by another process. Resident System Processes run code that is a part of the operating system itself. RSPs become an integral part of the operating system image during system generation. They are immediately available to perform operating system tasks. For example, the CLOCK process is an RSP that maintains the time of day within the operating system.

The following list briefly summarizes Concurrent CP/M's capabilities.

- Interprocess communication, synchronization, and mutual exclusion functions are provided by system queues.
- A logical interrupt mechanism using flags allows Concurrent CP/M to interface with any physical interrupt structure.
- System timing functions enable processes running under Concurrent CP/M to compute elapsed times, delay execution for specified intervals, and to access and set the current date and time.
- Shared file system allows multiple programs to access common data files while maintaining data integrity.
- Shared code support eliminates program loading of another copy of the same program and conserves memory space.
- 8087 support takes advantage of fast 8087 math instructions.
- Virtual console handling lets a single user run multiple programs, each in its own console environment.
- Real-time process control allows communications and data acquisition without loss of information.

Functionally, Concurrent CP/M is composed of several distinct modules, as shown in Figure 1-2.




Figure 1-2. Concurrent CP/M Functional Modules

- The Supervisor (SUP)
- The Real-time Monitor (RTM)
- The Memory Management Module (MEM)
- The Character I/O Module (CIO)
- The Virtual Console Screen Manager
- The Basic Disk Operating System (BDOS)
- The Extended I/O System (XIOS)
- The Terminal Message Processor (TMP)

The SUP module handles miscellaneous system calls such as returning the version number or the address of the System Data Area. SUP also calls other system calls when necessary.

The RTM module monitors the execution of running processes and arbitrates conflicts for the system's resources.

The MEM module allocates and frees memory upon demand from executing processes.

The CIO module handles all character I/O for console and list devices in the system.

The Virtual Console Screen Manager extends the CIO to support virtual console environments.

The BDOS is the hardware-independent module that contains the logically invariant portion of the file system for Concurrent CP/M. The BDOS file system is explained in detail in Section 2.

The XIOS is the hardware-dependent module that defines the interface of Concurrent CP/M to a specific hardware environment. See the *Concurrent CP/M System Guide* for an explanation of the XIOS.

When Concurrent CP/M is executing a single program on a single virtual console, its speed approximates that of CP/M-86. But when multiple processes are running on several virtual consoles, the execution of each individual process slows according to the proportion of I/O to CPU resources it requires. A process that performs a large amount of I/O in proportion to computing exhibits only minor speed degradation. This also applies to a process that performs a large amount of computing, but runs concurrently with other processes that are largely I/O-bound. On the other hand, significant speed degradation occurs where more than one compute-bound process is running.

1.2 Supervisor (SUP)

The Supervisor module (SUP) manages the interface between processes and the operating system kernel. It also manages internal communication between operating system modules. All system calls, whether they originate from a transient process or internally from another system module, go through a common table-driven function interface in SUP. SUP also handles the P_LOAD (Load Process) and P_CLI (Call Command Line Interpreter) system calls.

1.3 Real-time Monitor (RTM)

The Real-time Monitor (RTM) is the real-time multitasking nucleus of Concurrent CP/M. The RTM performs process dispatching, queue management, flag management, device polling, and system timing tasks. User programs can also call many of the RTM system calls used to perform these tasks.

1.3.1 Process Dispatching

Although Concurrent CP/M is a multiprocess operating system, only one process has access to the CPU resource at any given time. Unless you specifically write a program to communicate or synchronize execution with other processes, a process is unaware of other processes competing for system resources.

The primary task of the RTM is to transfer, or dispatch, the CPU resource from one process to another. The RTM module called the Dispatcher performs this task. The RTM maintains two data structures, the Process Descriptor (PD) and the User Data Area (UDA), for each process running under Concurrent CP/M. The Dispatcher uses these data structures to save and restore the current state of each running process.

Each process in the system resides in one of three states: ready, running, or suspended. A ready process is one that is waiting for the CPU resource only. A running process is one that the CPU is currently executing. A suspended process is one that is waiting for a system resource or a specified event, such as the occurrence of an interrupt, an indication that polled hardware is ready, or the expiration of a delay period.

Any existing process is represented on a system list. The Dispatcher removes a process from one list and places it on another. The Process Descriptor of the currently running process is the first entry on the Ready List. Other processes ready to run are represented on the Ready List in order of priority. Suspended processes are on other system lists, depending on why the processes were suspended.

A dispatch operation can be summarized as follows:

1. The Dispatcher suspends the process from execution and stores its current state in the Process Descriptor and the UDA.
2. The Dispatcher places the process on an appropriate system list, depending on why the Dispatcher was called. For example, if a process is to delay for a certain number of system ticks, its Process Descriptor is placed on the Delay List. When a process releases a resource, the process is usually placed back on the Ready List. If another process is waiting for the resource, that process is taken off its current system list and also placed on the Ready List.
3. The highest priority process on the Ready List is chosen for execution. If two or more processes have the same priority, the process that has waited the longest executes first.
4. The Dispatcher restores the state of the selected process from its Process Descriptor and UDA, and gives it the CPU resource.
5. The process executes until it needs a busy resource, a resource needed by another process becomes available, or an interrupt occurs. At this point, a dispatch occurs, allowing another process to run.

Only processes on the Ready List are eligible for selection during dispatch. By definition, a process is on the Ready List if it is waiting only for the CPU resource. Processes waiting for other system resources cannot execute until the resources they require are available. Concurrent CP/M blocks a process from execution if it is waiting for:

- a queue message so it can complete a Q_READ operation.
- space to become available in a queue so it can complete a Q_WRITE operation.
- a console or list device to become available.
- a specified number of system clock ticks before it can be removed from the system Delay List.
- an I/O event to complete.

These situations are discussed in greater detail in the following sections.

A running process not needing a resource and not releasing one runs until an interrupt causes a dispatch. While not all interrupts cause dispatches, the system clock generates interrupts every clock tick and forces a dispatch each time. Clock ticks usually occur 60 times a second (approximately every 16.67 milliseconds), and allow time sharing within a real-time environment.

Concurrent CP/M is a priority-driven system. This means that during a dispatch, the operating system gives the CPU resource to the process with the best priority. The Dispatcher allots equal shares of the system's resources to processes with the same priority. With priority dispatching, the system never passes control to a lower-priority process if there is a higher-priority process on the Ready List. Because high-priority, compute-bound processes tend to monopolize the CPU resource, it is best to reduce their priority to avoid degrading overall system performance.

1.3.2 Queue Management

Queues perform several critical functions for processes running under Concurrent CP/M. A process can use a queue for communicating with another process, synchronizing its execution with that of another process, and for exclusion of other processes from protected system resources. A process can make, open, delete, read from, or write to a queue with system calls similar to those used to manage disk files.

Each system queue consists of two parts: the queue descriptor, and the queue buffer. Concurrent CP/M implements these special data structures as memory files that contain room for a specified number of fixed-length messages.

When the `Q_MAKE` system call creates a queue, this queue is assigned a unique 8-character name. As the name queue implies, messages are read from a queue on a first-in, first-out basis.

A process can read from or write to a queue conditionally or unconditionally. If the queue is empty when a conditional read is performed, or full when a conditional write is performed, the system returns an error code to the calling process. On the other hand, if a process attempts an unconditional queue operation in these circumstances, the system suspends it from execution until the operation becomes possible.

More than one process can wait to read or write a queue message from the same queue at the same time. When these operations become possible, the system restores the highest priority process first; processes with the same priority are restored on a first-come, first-served basis.

Mutual exclusion queues are a special type of queue under Concurrent CP/M. They contain one message of zero length and their names follow a convention, beginning with the uppercase letters `MX`. A mutual exclusion queue acts as a binary semaphore, ensuring that only one process uses a resource at any time.

Access to a resource protected by a mutual exclusion queue takes place as follows:

1. A process issues an unconditional Q_READ call to the MX queue protecting the resource, thereby suspending itself if the message is not available.
2. When the message becomes available, the process accesses the protected resource. Note that from the time the process issues the unconditional read, any other process attempting to access the same resource is suspended.
3. The process writes the zero-length message back to the queue when it has finished using the protected resource, thus freeing the resource for other processes.

As an example, the system mutual exclusion queue, MXdisk, ensures that processes cannot access the file system simultaneously. Note that the BDOS, not the application software, executes the preceding series of queue calls. Therefore the mutual exclusion process is transparent to the programmer, who is only responsible for originating the disk system calls.

Mutual exclusion queues differ from normal queues in another way. When a process reads a message from a mutual exclusion queue, the RTM notes the Process Descriptor address within the Queue Descriptor. This establishes the owner of the queue message. If the operating system aborts the process while it owns the mutual exclusion message, the RTM automatically writes the message back to all mutual exclusion queues whose messages are owned by the aborted process. This grants other processes access to protected resources owned by the aborted process.

1.3.3 System Timing Functions

Concurrent CP/M's timing system calls include keeping the time of day and delaying the execution of a process for a specified period of time. An internal process called CLOCK provides the time of day for the system. This process issues DEV_WAITFLAG system calls on the system's one second flag, Flag 2. When the XIOS Tick Interrupt Handler sets this flag, it initiates the CLOCK process, which then increments the internal time and date.

Subsequently, the CLOCK process makes another DEV_WAITFLAG call and suspends itself until the flag is set again. Concurrent CP/M provides system calls that allow you to set and access the internal date and time. In addition, the file system uses the internal time and date to record when a file is updated, created, or last accessed.

The P_DELAY system call replaces the typical programmed delay loop for delaying process execution. P_DELAY requires that Flag 1, the system tick flag, be set approximately every 16.67 milliseconds, or 60 times a second; the XIOS Tick Interrupt Handler also sets this flag. When a process makes a P_DELAY system call, it specifies the number of ticks for which the operating system is to suspend it from execution. The system maintains the address of the Process Descriptor for the process on an internal Delay List along with its current delay tick count. When a DEV_SETFLAG call occurs, setting Flag 1, the tick count is decremented. When the delay count goes to zero, the system removes the process from the Delay List and places it on the Ready List.

Note: The length of a tick might vary from installation to installation. For instance, in Europe, a tick is commonly 20 milliseconds, yielding 50 ticks per second. The description of the P_DELAY system call in Section 6 describes how to determine the correct number of ticks to delay 1 second.

1.4 Memory Module (MEM)

Concurrent CP/M supports an extended, fixed partition model of memory management; the Memory Module handles all memory management system calls. In practice, the exact method that the operating system uses to allocate and free memory is transparent to the application program. Therefore you should take care to write code independent of the memory management model; use only the Concurrent CP/M specific memory system calls described in Section 6.

1.5 Basic Disk Operating System (BDOS)

Except for auxiliary device support, Concurrent CP/M BDOS is an upward-compatible version of the single-tasking CP/M-86 BDOS. It handles file creation and deletion, facilitates sequential or random file access, and allocates and frees disk space. In most cases, CP/M-86 programs that make BDOS calls for I/O can run under Concurrent CP/M without modification. Concurrent CP/M's BDOS is extended to provide support for multiple virtual consoles and list devices. In addition, the file system is extended to provide services required in a multitasking environment. The major extensions to the file system are

- File locking. Files opened under Concurrent CP/M cannot be opened or deleted by other tasks. This feature prevents accidental conflicts with other tasks.

- Shared access to files. As a special option, independent users can open the same file in shared or unlocked mode. Concurrent CP/M supports record locking and unlocking commands for files opened in this mode and protects files opened in shared mode from deletion by other tasks.
- Date Stamps. The BDOS optionally supports two time and date stamps, one recording when a file is updated, and the other recording when the file was created or last accessed.
- Password Protection. The password protection feature is optional at either the file or drive level. The operator or applications program assigns disk drive passwords, while application programs can assign file protection passwords in several modes.
- Extended Error Module. Besides the default error mode, Concurrent CP/M has two optional error-handling modes that return an error code to the calling process in the event of an unrecoverable disk error.

1.6 Character I/O Module (CIO)

The Character I/O module handles all console and list I/O. Under Concurrent CP/M, every character I/O device is associated with a data structure called a Console Control Block (CCB) or a List Control Block (LCB). These data structures reside in the XIOS. The CCB contains the current owner, status information, line editing variables, and the root of a linked list of Process Descriptors (PDs) that are waiting for access. More than one process can wait for access to a single console. These processes are maintained on a linked list of Process Descriptors in priority order. The LCBs contain similar information about the list devices. See the *Concurrent CP/M System Guide* for more information about LCBs and CCBs.

1.7 Virtual Console Screen Management

Virtual console screen management is coordinated by four separate modules: the CIO, the PIN (Physical INput) and VOUT (Virtual OUTput) processes, and the XIOS. The line editing associated with the C_READSTR call is performed in the CIO. The PIN process handles keyboard input for all the virtual consoles; it also traps and implements the CTRL-C, CTRL-S, CTRL-Q, CTRL-P, and CTRL-O functions. The VOUT process spools console output from processes running on background buffered mode consoles, and handshakes with the PIN process to display spooled console output when the background console is brought to the foreground. The XIOS decides which special keys represent the virtual consoles, and returns a special code from IO_CONIN when you request a screen switch. The XIOS also implements any screen saving and restoring when screens are switched. See the *Concurrent CP/M System Guide* and the discussion of the IO_SWITCH function.

The PIN process reads the keyboard by directly calling the XIOS IO_CONIN function. This is the only place in the operating system IO_CONIN is called. The PIN scans the input stream from the keyboard for switch screen requests and the special function keystrokes CTRL-C, CTRL-S, CTRL-Q, CTRL-P, and CTRL-O. All other keyboard input is written to the VINQ (Virtual Console INput Queue) associated with the foreground virtual console. The data in the VINQ becomes a type-ahead buffer for each virtual console, and is returned to the process attached to that console as it performs console input.

When PIN sees a CTRL-C it calls P_ABORT to abort the process attached to the virtual console, flushes the type-ahead buffer in the VINQ, turns off CTRL-S, and performs a DRV_RESET call for each logged-in drive. The P_ABORT call succeeds when the Process Keep flag is not on, saving the Terminal Message Processes (refer to P_CREATE for information on the process descriptor). The DRV_RESET calls affect only the removable media drives, as specified in the CKS field of the Disk Parameter Blocks in the XIOS (refer to the *Concurrent CP/M System Guide* for further details on Disk Parameter Blocks).

CTRL-S stops any output to the screen. CTRL-S stays set when a virtual console is switched to the background.

CTRL-O discards any console output to the virtual console. CTRL-O is turned off when any other key is subsequently pressed, except for the keys representing the virtual consoles.

CTRL-P echoes console output to the default list device specified in the LIST field of the process descriptor attached to the virtual console. If the list device is attached to a process, a PRINTER BUSY message appears.

All of the above control keys can be disabled by the C_MODE call. When one of the above control characters is disabled with C_MODE or when the process owning the virtual console is using the C_RAWIO call, the PIN does not act on the control character but instead writes it to the VINQ. It is thus possible to read any of the above control characters from an application program. These control keys are discussed in depth in the *Concurrent CP/M User's Guide*.

1.8 Extended Input/Output System (XIOS)

The XIOS module is similar to the CP/M-86 Basic Input/Output System (BIOS) module, but it is extended in several ways. Primitive operations, such as console I/O, are modified to support multiple virtual consoles. Several new primitive system calls, such as DEV_POLL, support Concurrent CP/M's additional features, including elimination of wait loops for real-time I/O operations.

1.9 Terminal Message Processes (TMP)

The Concurrent CP/M Terminal Message Processes (TMPs) are resident system processes that accept command lines from the virtual consoles and call the Command Line Interpreter (CLI) to execute them. The TMP prints the prompt on the virtual consoles.

Each virtual console has an independent TMP defining that console's environment, including default disk, user number, printer, and console.

1.10 Transient Programs

Under Concurrent CP/M, a transient program is one that is not system-resident. The system must load such programs from disk into available memory each time they execute. The command file of a transient program is identified by the filetype CMD. When you enter a command at the console, the operating system searches on disk for the appropriate CMD file, loads it, and initiates it. Concurrent CP/M supports three different execution models for transient programs: the 8080 Model, the Small Model, and the Compact Model. Sections 4.1.1 through 4.1.3 describe these models in detail.

1.11 System Call Calling Conventions

When a Concurrent CP/M process makes a system call, it loads values into the registers shown in Table 1-1 and initiates Interrupt 224 (via the INT 224 instruction), reserved by the Intel Corporation for this purpose.

Table 1-1. Registers Used by System Calls

ENTRY PARAMETERS	
Register	CL: System Call Number
	DL: Byte Parameter
	or
	DX: Word Parameter
	or
	DX: Address - Offset
	DS: Address - Segment
RETURN VALUES	
Register	AL: Byte Return
	or
	AX: Word Return
	or
	AX: Address - Offset
	ES: Address - Segment
	BX: Same as AX
	CX: Error Code

Concurrent CP/M preserves the contents of registers SI, DI, BP, SP, SS, DS, and CS through the operating system calls. The ES register is preserved when it is not used to hold a return segment value. Error codes returned in CX are shown in Table 6-5, CX Error Codes.

1.12 SYSTAT: System Status

The SYSTAT utility is a development tool that shows the internal state of Concurrent CP/M. SYSTAT describes memory allocation, current processes, system queue activity, and many informative parameters associated with these system data structures. Furthermore, SYSTAT presents two views: either a static snapshot of system activity, or a continuous, real-time window into Concurrent CP/M.

You can specify SYSTAT in one of two modes. If you know which display you want, you can specify it in the invocation, using an option shown in the menu below. If you do not specify an option, select a display from this menu by typing

```
A>SYSTAT <cr>
```

The screen clears and the main menu appears:

```
Which Option?
```

```
H(elp)
M(emory)
O(verview)
P(rocesses - All)
Q(ueues)
U(ser Processes)
C(onsoles)
E(xit)
->_
```

Press the appropriate letter to obtain a display.

When you select H(elp), the HELP file demonstrates the proper syntax and available options:

To use SYSTAT with the menu: At the system prompt type SYSTAT <CR>

To use SYSTAT without the menu: At the system prompt type the command

```
SYSTAT [option] -or-
SYSTAT [option C] -or-
SYSTAT [option C ###]
```

```
-where-
-> option =
    M(emory) P(rocesses) O(overview) C(onsoles)
    U(ser Processes) Q(ueues) H(elp)

-> C = Continuous display
    ## = 1-2 digits indicating the period,
        in seconds, between display refreshes.
```

Type any letter to return to the menu.

The M, P, Q, and U and C options ask you if you prefer a continuous display. If you type y, Concurrent CP/M asks for a time interval, in seconds, and then displays a real-time window of information. If you type n, a static snapshot of the requested information appears. In either case, press any key to return to the menu.

The M(emory) option displays all memory potentially available to you, but it does not display restricted memory. The partitions are listed in memory-address order. Length parameter is shown in paragraph values.

The O(overview) option displays an overview of the system parameters, as specified at system generation time. The display is not continuous.

The P(rocess) option displays all system processes and the resources they are using.

The Q(ueues) option displays all system queues, listing queue readers, writers, and owners.

The U(ser Processes) option displays only user-initiated processes in the same format as the P(rocess) option.

The C(onsoles) option displays console information; that is, background, foreground, buffered, suspended, purging, CTRL-Q, and so on.

The E(xit) option returns you to system level from the menu, as does CTRL-C.

End of Section 1

Section 2

The Concurrent CP/M File System

2.1 File System Overview

The Basic Disk Operating System (BDOS) file system supports from one to sixteen logical drives. Each logical drive has two regions: a directory area and a data area. The directory area defines the files that exist on the drive and identifies the data area space that belongs to each file. The data area contains the file data defined by the directory.

The directory area consists of sixteen logically independent directories. These directories are identified by user numbers 0 through 15. During execution, a process runs with a system parameter called the user number set to a single value. The user number specifies the current active directories for all drives on the system. For example, the Concurrent CP/M DIR utility displays only files within a directory selected by the current user number.

The file system automatically allocates directory and data area space when a process creates or extends a file, and returns previously allocated space to free space when a process deletes or truncates a file. If no directory or data space is available for a requested operation, the BDOS returns an error code to the calling process. The allocation and retrieval of directory and data space is transparent to the calling process. As a result, you need not be concerned with directory and drive organization when using the file system calls.

An eight-character filename and a three-character filetype field identify each file in a directory. Together, these fields must be unique for each file within a directory. However, files with the same filename and filetype can reside in different user directories without conflict. Processes can also assign an eight-character password to a file to protect it from unauthorized access.

All system calls that involve file operations specify the requested file by filename and filetype. For some system calls, multiple files can be specified by a technique called ambiguous reference. This technique uses question marks and asterisks as wildcard characters to give the file system a pattern to match as it searches a directory.

The file system supports two categories of system calls: file-access system calls and drive-related system calls. The file-access system calls have mnemonics beginning with F_, and the drive-related system calls have mnemonics beginning with DRV_. The next two sections introduce the file system calls.

2.1.1 File-access System Calls

Most of the file-access system calls can be divided into two groups: system calls that operate on files within a directory and system calls that operate on records within a file. However, the file-access category also includes several miscellaneous functions that either affect the execution of other file-access system calls or are commonly used with them.

System calls in the first file-access group include calls to search for one or more files, delete one or more files, rename or truncate a file, set file attributes, assign a password to a file, and compute the size of a file. Also included in this group are system calls to open a file, to create a file, and to close a file.

The second file-access group includes system calls to read or write records to a file, either sequentially or randomly, by record position. BDOS read and write system calls transfer data in 128-byte units, which is the basic record size of the file system. This group also includes system calls to lock and unlock records and thereby allows multiple processes to have coordinated access to records within a commonly accessed file.

Before making read, write, lock, or unlock system calls for a file, you must first open or create the file. Creating a file has the side effect of opening the file for record access. In addition, because Concurrent CP/M supports three different modes of opening files (Locked, Unlocked, and Read-Only), there can be other restrictions on system calls in this group that are related to the open mode. For example, you cannot write to a file that you have opened in Read-Only mode.

After a process has opened a file, access to the file by other processes is restricted until the file is closed. Again, the exact nature of the restrictions depends on the open mode. However, in all cases the file system does not allow a process to delete, rename, or change a file's attributes if another process has opened the file. Thus, the F_CLOSE system call performs two steps to terminate record access to a file. It permanently records the current status of the file in the directory and removes the open-file restrictions limiting access to the file by other processes.

The miscellaneous file-access system calls include calls to set the current user number, set the DMA address, parse an ASCII file specification and set a default password. This group also includes system calls to set the BDOS Multisector Count and the BDOS Error Mode. The BDOS Multisector count determines the number of 128-byte records to be processed by the read, write, lock, and unlock system calls. The Multisector count can range from 1 to 128; the default value is one. The BDOS Error Mode determines whether the file system intercepts certain errors or returns on all errors to the calling process.

2.1.2 Drive-related System Calls

BDOS drive-related system calls select the default drive, compute a drive's free space, interrogate drive status, and assign a directory label to a drive. A drive's directory label controls whether the file system enforces file password protection for files in the directory. It also specifies whether the file system is to perform date and time stamping of files on the drive.

This category also includes system calls to reset specified drives and to control whether other processes can reset particular drives. When a drive is reset, the next operation on the drive reactivates it by logging it in. Logging in a drive initializes the drive for directory and file operations. The purpose of a drive reset call is to prepare for a media change on drives that support removable media. Under Concurrent CP/M, drive reset calls are conditional. A process cannot reset a drive if another process has a file open on the drive.

The following table summarizes the BDOS file system calls.

Table 2-1. File System Calls

<i>Mnemonic</i>	<i>Description</i>
DRV_ACCESS	Access Drive
DRV_ALLOCVEC	Get Drive Allocation Vector
DRV_ALLRESET	Reset All Drives
DRV_DPB	Get Disk Parameter Block Address
DRV_GET	Get Default Drive
DRV_GETLABEL	Get Directory Label
DRV_FLUSH	Flush Data Buffers
DRV_FREE	Free Drive
DRV_LOGINVEC	Return Logged In Vector
DRV_RESET	Reset Drive
DRV_ROVEC	Return R/O Vector
DRV_SETLABEL	Set Directory Label

Table 2-1. (continued)

<i>Mnemonic</i>	<i>Description</i>
DRV_SET	Set (Select) Drive
DRV_SETRO	Set Drive To Read-Only
DRV_SPACE	Get Free Space On Drive
F_ATTRIB	Set File's Attributes
F_CLOSE	Close File
F_DELETE	Delete File
F_DMASEG	Set DMA Segment
F_DMAGET	Get DMA Address
F_DMAOFF	Set DMA Offset
F_ERRMODE	Set BDOS Error Mode
F_LOCK	Lock Record In File
F_MAKE	Make A New File
F_MULTISEC	Set BDOS Multisector Count
F_OPEN	Open File
F_PARSE	Parse Filename
F_PASSWD	Set Default Password
F_RANDREC	Return Record Number For File Read-Write
F_READ	Read Record Sequentially From File
F_READRAND	Read Random Record From File
F_RENAME	Rename File
F_SIZE	Compute File Size
F_SFIRST	Directory Search First
F_SNEXT	Directory Search Next
F_TIMEDATE	Return File Time/Date Stamps Password Mode
F_TRUNCATE	Truncate File
F_UNLOCK	Unlock Record In File
F_USERNUM	Set/Get Directory User Number
F_WRITE	Write Record Sequentially Into File
F_WRITERAND	Write Random Record Into File
F_WRITEXFCB	Write File's XFCB
F_WRITEZF	Write Random Record With Zero Fill

The following sections contain information on important topics related to the file system. Read these sections carefully before attempting to use the system calls described individually in Section 6.

2.2 File Naming Conventions

Under Concurrent CP/M, a file specification consists of four parts: a drive specifier, the filename field, the filetype field, and the file password field. The general format for a command line file specification is shown below:

```
{d:} filename { .typ } {;password}
```

The drive specifier field specifies the drive where the file is located. The filename and filetype fields identify the file. The password field specifies the password if a file is password protected.

The drive, type, and password fields are optional, and delimiters are required only when specifying their associated fields. The drive specifier can be assigned a letter from A to P, where the actual drive letters supported on a given system are determined by the XIOS implementation. When the drive letter is not specified, the current default drive is assumed.

The filename and password fields can contain one to eight non-delimiter characters. The filetype field can contain one to three non-delimiter characters. All three fields are left justified and padded with blanks, if necessary. Omitting the optional type or password fields implies a field specification of all blanks.

Under Concurrent CP/M, the P_CLI system call interprets ASCII command lines and loads programs. The P_CLI system call makes F_PARSE system calls to parse file specifications from a command line. F_PARSE recognizes certain ASCII characters as delimiters when it parses a file specification. These characters are shown in Table 2-2.

Table 2-2. Valid Filename Delimiters

<i>ASCII</i>	<i>Hex Equivalent</i>
null	000H
space	020H
return	00DH
tab	009H
:	03AH
.	02EH
;	03BH
=	03DH
,	02CH
[05BH
]	05DH
<	03CH
>	03EH
	07CH

The F_PARSE system call also excludes all control characters from the file specification fields and translates all lowercase letters to uppercase.

Avoid using parentheses and the backslash character, \, in the filename and filetype fields because they are commonly used delimiters. Use asterisk and question mark characters, * and ?, only to make an ambiguous file reference. When F_PARSE encounters an asterisk in a filename or filetype field, it pads the remainder of the field with question marks. For example, a filename of X*.* is parsed to X??????.???. The BDOS F_SFIRST, F_SNEXT, and F_DELETE system calls match a question mark in the filename or filetype fields to the corresponding position of any directory entry belonging to the current user number. Thus, a search operation for X??????.??? finds all the files in the current user directory beginning in X. Most other file-access BDOS system calls treat the presence of a question mark in the filename or filetype fields as an error.

It is not mandatory to follow the file naming conventions of Concurrent CP/M when you create or rename a file with BDOS system calls directly from an application program. However, the conventions must be used if the file is to be accessed from a command line. For example, the P_CLI system call cannot locate a command file in the directory if its filename or filetype field contains a lowercase letter.

As a general rule, the filetype field names the generic category of a particular file, and the filename field distinguishes individual files within each category. Although they are generally arbitrary, Table 2-3 lists some of the generic filetype categories that have been established.

Table 2-3. Filetype Conventions

<i>Filetype</i>	<i>Description</i>
A86	8086 Assembler Source
ASM	8080 Assembler Source
BAK	Text or Source Back-up
BAS	BASIC Source File
C	C Source File
CMD	8086 Command File
COM	8080 Command File
CON	CCP/M Modules
DAT	Data File
HEX	ASM80 HEX File
H86	ASM86 HEX File
INT	Intermediate File
LIB	Library File
L86	Library File
LST	List File
PLI	PL/I Source File
PRL	Page Relocatable
REL	Relocatable Module
RSP	Resident System Process
SPR	System Page Relocatable
SUB	SUBMIT File
SYM	Symbol File
SYS	System File
\$\$\$	Temporary File

2.3 Disk Drive and File Organization

The file system can support up to sixteen logical drives, identified by the letters A through P. A logical drive usually corresponds to a physical drive on the system, particularly for physical drives that support removable media such as floppy disks. High-capacity hard disks, however, are commonly divided into multiple logical drives. If a disk contains system tracks reserved for the boot loader, these tracks precede the tracks of the disk mapped by the logical drive. In this manual, references to drives mean logical drives, unless explicitly stated otherwise.

The maximum file size supported on a drive is 32 megabytes. The maximum capacity of a drive is determined by the data block size specified for the drive in the XIOS. The data block size is the basic unit in which the BDOS allocates space to files. Table 2-4 displays the relationship between data block size and total drive capacity.

Table 2-4. Drive Capacity

<i>Data Block Size</i>	<i>Maximum Drive Capacity</i>
1K	256 kilobytes
2K	64 megabytes
4K	128 megabytes
8K	256 megabytes
16K	512 megabytes

Each drive is divided into two regions: a directory area and a data area. The directory area contains from one to sixteen blocks located at the beginning of the drive. The actual number is set in the XIOS. Directory entries residing in this area define the files that exist on the drive. In addition, the directory entries belonging to a file identify the data blocks in the drive's data area that contain the file's records. The directory area is logically subdivided into sixteen independent directories identified as user 0 through 15. Each independent directory shares the actual directory area on the drive.

Each disk file may consist of a set of up to 262,144 (40000H) 128-byte records. Each record of a file is identified by its position in the file. This position is called the record's Random Record Number. If a file is created sequentially, the first record has a position of zero, while the last record has a position one less than the number of records in the file. Such a file can be read sequentially, beginning at record zero, or randomly by record position. Conversely, if a file is created randomly, records are added to the file by specified position. A file created in this way is called sparse if positions exist within the file where a record has not been written.

The BDOS automatically allocates data blocks to a file to contain the file's records on the basis of the record positions consumed. Thus, a sparse file that contains two records, one at position zero, the other at position 262,143, consumes only two data blocks in the data area. Sparse files can be created and accessed only randomly, not sequentially. Note that any data block allocated to a file is permanently allocated until the file is deleted or truncated. These are the only mechanisms supported by the BDOS for releasing data blocks belonging to a file.

Source files under Concurrent CP/M are treated as a sequence of ASCII characters, where each line of the source file is followed by a carriage return/line-feed sequence, 0DH followed by 0AH. Thus, a single 128-byte record could contain several lines of source text. The end of an ASCII file is denoted by a CTRL-Z character (1AH), or a real end-of-file, returned by the BDOS read system call. Note that these source file conventions are not supported in the file system directly but are followed by Concurrent CP/M utilities such as TYPE and ASM-86[™]. In addition, CTRL-Z characters embedded within other types of files such as CMD files do not signal end-of-file.

2.4 File Control Block Definition

The File Control Block (FCB) is a system data structure that serves as an important channel for information exchange between a process and BDOS file-access system calls. A process initializes an FCB to specify the drive location, filename and filetype fields, and other information that is required to make a file-access call. For example, in an F_OPEN system call, the FCB specifies the name and location of the file to be opened. In addition, the file system uses the FCB to maintain the current state and record position of an open file. Some file-access system calls use special fields within the FCB for invoking options. Other file-access system calls use the FCB to return data to the calling program. All BDOS random I/O system calls require the calling process to specify the Random Record Number in a 3-byte field at the end of the FCB.

When a process makes a BDOS file-access system call, it passes an FCB address to the BDOS. This address has two 16-bit components: register DX, which contains the offset, and register DS, which contains the segment. The length of the FCB data area depends on the BDOS system call. For most system calls, the minimum length is 33 bytes. For the F_READRAND, F_WRITERAND, F_WRITEZF, F_LOCK, F_UNLOCK, F_RAND REC, F_SIZE, and F_TRUNCATE system calls, the minimum FCB length is 36 bytes. When the F_OPEN or F_MAKE system calls open a file in Unlocked mode, the FCB must be at least 35 bytes long. Figure 2-1 displays the FCB data structure in two formats.




Figure 2-1. FCB - File Control Block

The fields in the FCB are defined as follows:

Table 2-5. FCB Field Definitions

<i>Field</i>	<i>Definitions</i>
DR	<p>Drive Code (0-16).</p> <p>0 => use default drive for file 1 => auto disk select drive A 2 => auto disk select drive B</p> <p style="text-align: center;">.</p> <p>16=> auto disk select drive P</p>
F1...F8	<p>Contain the filename in ASCII uppercase, with high bit = 0. F1', ..., F8' denote the high-order bit of these positions and are called attribute bits.</p>
T1,T2,T3	<p>Contain the filetype in ASCII uppercase, with high bit = 0. T1', T2', and T3' denote the high bit of these positions and are also called attribute bits.</p> <p>T1' = 1 => Read-Only file, T2' = 1 => System file, T3' = 1 => File has been archived.</p>
EX	<p>Contains the current extent number. This field is initialized to 0 by the calling process, but it can range from 0 to 31 during file I/O.</p>
CS	<p>Contains the FCB checksum value for open FCBs.</p>
RS	<p>Reserved for internal system use</p>
RC	<p>Record count for extent EX. This field takes on values from 0 to 255 (values greater than 128 imply a record count of 128).</p>

Table 2-5. (continued)

<i>Field</i>	<i>Definitions</i>
D0...D15	Normally filled in by Concurrent CP/M and reserved for system use. Also used to specify the new filename and filetype with the F_RENAME system call.
CR	Current record to read or write in a sequential file operation. This field is normally set to zero by the calling process when a file is opened or created.
R0,R1,R2	Optional Random Record Number in the range 0-262,143 (0 - 3FFFFH). R0, R1, R2 constitute an 18-bit value with low byte R0, middle byte R1, and high byte R2.

Note: The 2-byte File ID is returned in bytes R0 and R1 of the FCB when a file is successfully opened in Unlocked mode (refer to Section 2.10).

2.4.1 FCB Initialization and Usage

The calling process must initialize bytes 0 through 11 of the referenced FCB before making the following file-access system calls: F_ATTRIB, F_DELETE, F_MAKE, F_OPEN, F_RENAME, F_SFIRST, F_SIZE, F_SNEXT, F_TIMEDATE, F_TRUNCATE, and F_WRITEXFCB. Normally, the DR field specified the drive location of the file, and the name and type fields specify the name of the file. You must also set the EX field of the FCB before calling F_MAKE, F_OPEN, F_SFIRST, and F_WRITEXFCB. Except for the F_WRITEXFCB system call, you can usually set this field to zero. Note that the F_RENAME system call requires the calling process to place the new filename and filetype in bytes D1 through D11.

The remaining file-access calls that use FCBs require an FCB that has been initialized by a prior file-access system call. For example, the F_SNEXT system call expects an FCB initialized by a prior F_SFIRST call. In addition, the F_LOCK, F_READ, F_READ-RAND, F_UNLOCK, F_WRITERAND, and F_WRITEZF system calls require an FCB that has been activated for record operations. Under Concurrent CP/M, only the F_OPEN and F_MAKE system calls can activate an FCB.

If you intend to process a file sequentially from the beginning, using the F_READ and F_WRITE system calls, you must set the CR field to zero before you make your first read or write call. In addition, when you make an F_LOCK, F_READRAND, F_UNLOCK, F_WRITERAND, or F_WRITEZF system call, you must set bytes R0 through R2 of the FCB to the requested Random Record Number. The F_TRUNCATE system call also requires the FCB random record field to be initialized.

The F_SFIRST, F_SNEXT, and F_DELETE system calls support multiple or ambiguous reference. In general, a question mark in the filename, filetype, or EX fields matches all values in the corresponding positions of directory entries during a directory search operation. File directory entries maintained in the directory area of each disk drive have the same format as FCBs except for byte 0, which contains the file's user number, and bytes 32 through 35, which are not present. The search system calls, F_SFIRST and F_SNEXT, also recognize a question mark in the FCB DR field, and, if specified, they return all directory entries on the disk regardless of user number, including empty entries. A directory FCB that begins with E5H is an empty or erased directory entry.

When the F_OPEN and F_MAKE system calls activate an FCB for record operations, they copy the FCB's matching directory entry from disk, excluding byte 0, into the FCB in memory. In addition, these system calls compute and store a checksum value in the CS field of the FCB. During subsequent record operations on the file, the file system uses this checksum field to verify that the FCB has not been modified by the calling process in an illegal way. Thus, all read, write, lock, and unlock operations on a file must specify a valid activated FCB; otherwise, the BDOS returns a checksum error. The BDOS performs this checking to protect the integrity of the file system. In general, you should not modify bytes 0 through 31 of an open FCB, except to set interface attributes (see Section 2.4.3). Other restrictions related to activated FCBs are discussed in Section 2.10.

The BDOS updates the memory copy of the FCB during file processing to maintain the current position within the file. During file write operations, the BDOS also updates the memory copy of the FCB to record the allocation of data blocks to the file. At the termination of file processing, the F_CLOSE system call permanently records this information on disk.

Note that the BDOS does not record the data blocks allocated to a file during write operations in the disk directory until the calling process issues an F_CLOSE call. Therefore, a process that creates or modifies files must close the files at the termination of file processing. Otherwise, data might be lost.

2.4.2 File Attributes

The high-order bits of the FCB filename (F1',...,F8') and filetype fields (T1',T2',T3') are called attribute bits. Attribute bits are 1-bit Boolean fields, where 1 indicates on or true, and 0 indicates off or false. Attribute bits indicate two kinds of attributes within the file system: file attributes and interface attributes. The file attributes are described in this section. Section 2.4.3 describes interface attributes.

The file attribute bits, F1',...,F4' and T1', T2', T3', indicate that a file has a defined attribute. These bits are recorded in a file's directory FCBs. File attributes can be set or reset only by the F_ATTRIB system call. When the F_MAKE system call creates a file, it initializes all file attributes to zero. A process can interrogate file attributes in an FCB activated by the F_OPEN system call, or in directory FCBs returned by the F_SFIRST and F_SNEXT system calls.

Note: The file system ignores the file attribute bits when it attempts to locate a file in the directory.

The file system defines file attributes T1',T2',and T3' as follows:

Table 2-6. File Attribute Definitions

<i>Attribute</i>	<i>Definition</i>
T1': Read-Only Attribute	This attribute, if set, prevents write operations to a file.
T2': System Attribute	This attribute, if set, identifies the file as a Concurrent CP/M system file. The Concurrent CP/M DIR utility does not usually display System files. In addition, user-zero system files can be accessed on a Read-Only basis from other user numbers.
T3': Archive Attribute	User-written archive programs use this attribute. When an archive program copies a file to back-up storage, it sets the archive attribute of the copied files. The file system automatically resets the archive attribute of a directory entry when writing to the directory entry's region of a file. An archive program can test this attribute in each of the file's directory entries using the F_SFIRST and F_SNEXT system calls. If all directory entries have the archive attribute set, the file has not been modified since the previous archive. The Concurrent CP/M PIP utility supports file archiving.

File attributes F1' through F4' of command files are defined as Compatibility Attributes under Concurrent CP/M (see Section 2.12). However, for all other files, attributes F1' through F4' are available for definition by the user.

2.4.3 Interface Attributes

The interface attributes are F5', F6', F7', and F8'. These attributes cannot be used as file attributes. Interface attributes F5' and F6' request options for BDOS file-access system calls. Table 2-7 lists the F5' and F6' attribute definitions for the system calls that define interface attributes. Note that the F5' = 0 and F6' = 0 definitions are not listed if their definition simply implies the absence of the associated option.

Table 2-7. BDOS Interface Attributes F5' and F6'

<i>System Call</i>	<i>Attribute</i>
F_ATTRIB	F5' = 1 : Maintain extended file lock F6' = 1 : Set file byte count
F_CLOSE	F5' = 1 : Partial Close F6' = 1 : Extend file lock
F_DELETE	F5' = 1 : Delete file XFCBs only and maintain extended file lock
F_LOCK	F5' = 0 : Exclusive Lock F5' = 1 : Shared Lock F6' = 0 : Lock existing records only F6' = 1 : Lock logical records
F_MAKE	F5' = 0 : Open in Locked mode F5' = 1 : Open in Unlocked mode F6' = 1 : Assign password to file
F_OPEN	F5' = 0 : Open in Locked mode F5' = 1 : Open in Unlocked mode F6' = 0 : Open in mode specified by F5' F6' = 1 : Open in Read-Only mode
F_RENAME	F5' = 1 : Maintain extended file lock
F_TRUNCATE	F5' = 1 : Maintain extended file lock
F_UNLOCK	F5' = 1 : Unlock all locked records

Section 6 details the above interface attribute definitions for each of the preceding system calls. Note that the BDOS always resets interface attributes F5' and F6' before returning to the calling process. Interface attributes F7' and F8' are reserved for internal use by the file system.

2.5 User Number Conventions

The Concurrent CP/M user facility divides each drive directory into sixteen logically independent directories, designated as user 0 through user 15. Physically, all user directories share the directory area of a drive. In most other aspects, however, they are independent. For example, files with the same name can exist on different user numbers of the same drive with no conflict. However, a single file cannot extend across more than one user number.

Only one user number is active for a specific process at one time. For this process, the current user number applies to all drives on the system. Furthermore, the FCB format does not contain a field that can override the current user number. As a result, all file and directory operations reference only directory entries associated with the current user number.

However, it is possible for a process to access files on different user numbers by setting the user number to the file's user number with the F_USERNUM system call before issuing the BDOS call. However, if a process attempts to read or write to a file under a user number different from the user number that was active when the file was opened, the file system returns an FCB checksum error.

When the P_CLI system call initiates a transient process or Resident System Process (described in detail in Section 5), it sets the user number to the default value established by the process issuing the P_CLI system call. The sending process is usually the TMP. However, the sending process can be another process, such as a transient program that makes a P_CHAIN call. A transient process can change its user number by making an F_USERNUM call. Changing the user number in this way does not affect the command line user number displayed by the TMP. Thus, when a transient process that has changed its user number terminates, the TMP restores and displays the original user number in the command line prompt when it regains control.

User 0 has special properties under Concurrent CP/M. The file system automatically opens files listed under user zero but requested under another user number if the file is not present under the current user number, and if the file on user zero has the system attribute (T2') set. This convention allows utilities, including overlays and any other commonly accessed files, to reside on user zero, but remain available to other users. This eliminates the need to copy commonly used utilities to all user numbers on a directory, and gives the Concurrent CP/M manager control over which files are directly accessible to the different user areas.

2.6 Directory Labels and XFCBs

The file system includes three special types of FCBs: the directory label and the XFCB, described in this section, and the SFCB, described in detail in Section 2.8.

The directory label specifies for its drive whether password support is to be activated, and if date and time stamping for files is to be performed. The format of the directory label is shown below in Figure 2-2.

DR	Name	Type	DL	S1	S2	RC	Password	TS1	TS2
00	01...	09...	12	13	14	15	16...	25..	29..

Figure 2-2. Directory Label Format

Table 2-8. Directory Label Field Definitions

<i>Field</i>	<i>Definition</i>
DR	drive code (0–16)
Name	directory label name
Type	directory label type
DL	directory label data byte Bit 7 - enable password support Bit 6 - perform access time stamping Bit 5 - perform update time stamping Bit 4 - perform create time stamping Bit 0 - Directory Label exists (Bit references are right to left, relative to 0)
S1,S2,RC	reserved for system use
Password	8-byte password field (encrypted)
TS1	4-byte creation time stamp field
TS2	4-byte update time stamp field

Only one directory label can exist in a drive's directory area. The directory label name and type fields are not used to search for a directory label; they can be used to identify a disk.

You can use the `DRV_SETLABEL` system call to create a directory label or update its fields. This system call can also assign a password to a directory label. The directory label password, if assigned, cannot be circumvented, whereas file password protection on a drive is an option controlled by the directory label. Thus, access to the directory label password provides the ability to bypass password protection on the drive.

Note: The file system provides no specific system call to read the directory label FCB directly. However, you can read the directory label data byte directly with the BDOS system call, DRV_GETLABEL. In addition, you can use the BDOS search system calls F_SFIRST and F_SNEXT to find a directory label. You can identify the directory label by a value of 32 (020H) in byte 0 of the directory FCB.

The XFCB is an extended FCB that can optionally be associated with a file in the directory. If present, it contains the file's password and password mode. The format of the XFCB is shown below in Figure 2-3.




Figure 2-3. XFCB - Extended File Control Block

The fields in the XFCB are defined in Table 2-9:

Table 2-9. XFCB Field Definitions

<i>Field</i>	<i>Definition</i>
DR	drive code (0–16)
File	filename field
Type	filetype field
PM	password mode Bit 7 - Read mode Bit 6 - Write mode Bit 5 - Delete mode (Bit references are right to left, relative to 0)
S1,S2,RC	reserved for system use
Password	8-byte password field (encrypted)
Reserved	8-byte area reserved for future use

An XFCB can be created only on a drive that has a directory label, and only if the directory label enables password protection. For drives in this state, there are two ways to create an XFCB for a file: with the F_MAKE system call or the F_WRITEXFCB system call. The F_MAKE system call creates an XFCB if the calling process requests that a password be assigned to the created file. The F_WRITEXFCB system call creates an XFCB when it is called to assign a password to an existing file. You can identify an XFCB in the directory by a value of 16 (010H) + N in byte 0 of the FCB, where N equals the user number.

2.7 File Passwords

There are two ways to assign passwords to a file: by the F_MAKE system call or by the F_WRITEXFCB system call. You can also change a file's password or password mode with the F_WRITEXFCB system call if you can supply the original password. Note that you cannot change a file's password or password mode if password protection for the drive is disabled by the directory label. However, even if you cannot supply a file's password, you can delete a file's XFCB, thereby removing its password protection, if password protection is disabled on the drive.

The Concurrent CP/M BDOS provides password protection in one of three modes when password support is enable by the directory label. Table 2-10 shows the difference in access level allowed to BDOS system calls when the password is not supplied.

Table 2-10. Password Protection Modes

<i>Mode</i>	<i>Access Level Allowed Without Password</i>
(1) Read	Cannot be read, modified, or deleted.
(2) Write	Can be read, but not modified or deleted.
(3) Delete	Can be read and modified, but not deleted.

If a file is password protected in Read mode, a process must supply the password to open the file. Processes cannot write to a file protected in Write mode without the password. A file protected in Delete mode allows read and write access, but a process must specify the password to delete or truncate the file, rename the file, or to modify the file's attributes. Thus, password protection in mode 1 implies mode 2 and 3 protection, and mode 2 protection implies mode 3 protection. All three modes require the user to specify the password to delete or truncate the file, rename the file, or to modify the file's attributes.

If a process supplies the correct password or the directory label disables password protection, then access to the BDOS system calls is the same as for a file that is not password-protected. In addition, the F_SFIRST and F_SNEXT system calls are not affected by file passwords. The following BDOS system calls test for passwords.

DRV_SETLABEL
F_ATTRIB
F_DELETE
F_OPEN
F_RENAME
F_WRITEXFCB
F_TRUNCATE

The BDOS maintains file passwords in the XFCB and directory label in encrypted form. To make a BDOS system call for a file that requires a password, a process must place the password in the first eight bytes of the current DMA, or make it the default password with the F_PASSWD system call, before making the system call.

Note: The BDOS maintains the assigned default password for each process. Processes inherit the default password of their parent process. You can set a given TMP's default password using the SET command; all programs loaded by this TMP inherit the same default password.

2.8 File Date and Time Stamps: SFCBs

The Concurrent CP/M file system uses a special type of directory entry called an SFCB to record date and time stamps for files. When a directory has been initialized for date and time stamping, SFCBs reside in every fourth position of the directory. Each SFCB maintains the date and time stamps for the previous three directory entries, as shown in Figure 2-4.




Figure 2-4. Directory Record with SFCB

This figure shows a 128-byte directory record containing an SFCB. Directory records have four directory entries, each 32 bytes long; SFCBs always occupy the last 32-byte entry in the directory record.

The SFCB itself contains five fields. The first field is a single byte containing the value 021H; this field identifies the SFCB within the directory. The next three fields, called the SFCB subfields, are each 10 bytes in length and contain the date and time stamps for their corresponding FCB entries in the directory record. The last byte of the SFCB is reserved for system use. Figure 2-5 shows the detail of the SFCB subfields.



Figure 2-5. SFCB Subfields

An SFCB subfield only contains valid information if its corresponding FCB in the directory record is an extent zero FCB. This FCB is a file's first directory entry. For password protected files, the SFCB subfield also contains the password mode of the file; the password mode field is zero for files without password protection. You can read SFCBs by making F_SFIRST and F_SNEXT system calls. In addition, you can make an F_TIMEDATE system call to retrieve the date and time stamps and password mode of a specified file. Refer to the T_GET system call definition in Section 6 for the description of the format of a date and time stamp field.

Concurrent CP/M supports three kinds of file stamping: create, access, and update. Create stamps record when the file was created, access stamps record when the file was last opened, and update stamps record the last time the file was modified. Create and access stamps share the same field. As a result, file access stamps overwrite any create stamps.

The directory label of a properly initialized disk determines the type of date and time stamping for files on the drive. The INITDIR utility initializes a directory for date and time stamping by placing an SFCB in every fourth directory entry. Disks not initialized in this way cannot support date and time stamping. In addition, date and time stamping is not performed if the disk's directory label is absent or does not specify date and time stamping, or if the disk is Read-Only.

Note that the directory label is also time stamped, but these stamps are not made in an SFCB; time stamp fields in the last eight bytes of the directory label show when it was created and last updated. Access stamping is not supported for directory labels.

The BDOS file system uses the system date and time when it records a date and time stamp. This value is maintained in a field in the SYSDAT part of the System Data Segment. The DATE utility sets the system time and date (refer to the *Concurrent CP/M User's Guide* for details of using DATE).

2.9 File Open Modes

The file system provides three different modes for opening files. They are defined below.

Locked Mode

A process can open a file in Locked mode only if the file is not currently opened by another process and the file is not a Read-Only file (attribute T1' set). Once open in Locked mode, no other process can open the file until it is closed. Thus, if a process successfully opens a file in Locked mode, that process owns the file until the file is closed or the process terminates. Files opened in Locked mode support read and write operations unless the file is password-protected in Write mode, and the process issuing the F_OPEN call cannot supply the password. In this case the BDOS allows only read operations to the file.

If a file opened in Locked mode is a Read-Only file, the F_OPEN system call automatically changes the open mode to Read-Only mode. Read-Only mode is described below.

Note: Locked mode is the Default mode for opening files under Concurrent CP/M.

Unlocked Mode

A process can open a file in Unlocked mode if the file is not currently open, or if another process has already opened the file in Unlocked mode. This mode allows more than one process to open the same file. Files opened in Unlocked mode support read and write operations unless the file is a Read-Only file (attribute T1' set) or the file is password-protected in Write mode and the process issuing the F_OPEN call cannot supply the password.

When opening a file in Unlocked mode, a process must reserve 35 bytes in the FCB because the F_OPEN system call returns a 2-byte value called the File ID in the R0 and R1 bytes of the FCB. The File ID is a required parameter for the F_LOCK and F_UNLOCK system calls. These BDOS system calls work only for files opened in Unlocked mode.

Read-Only Mode

A process can open a file in Read-Only mode if the file is not currently opened by another process or if another process has opened the file in Read-Only mode. This mode allows more than one process to open the same file for Read-Only access.

The F_OPEN system call performs the following steps for files opened in Locked or Read-Only mode. If the current user number is nonzero, and the file to be opened does not exist under the current user number, the F_OPEN system call searches the user zero directory for the file. If the file exists under user zero and has the system attribute T2' set, the BDOS opens the file under user zero. The open mode is automatically forced to Read-Only when this is done.

The F_OPEN and F_MAKE system calls use FCB interface attributes F5' and F6' to specify the open mode. The interface attribute definitions for these functions are listed in Table 2-7.

Note: The F_MAKE system call does not allow opening the file in Read-Only mode.

2.10 File Security

In general, the security measures implemented in the file system prevent accidental collisions between running processes. It is not possible to provide total security under Concurrent CP/M because the file system maintains file allocation information in open FCBs in the user's memory region, and Concurrent CP/M does not require memory protection. However, the file system is designed to ensure that multiple processes can share the same file system without interfering with each other by

- performing checksum verification of open FCBs.
- monitoring all open files and locked records via the system Lock List.

The BDOS validates the checksum of user FCBs before all I/O operations to protect the integrity of the file system from corrupted FCBs. The F_OPEN and F_MAKE system calls compute and assign checksums to FCBs. The F_READRAND, F_READ, F_WRITERAND, F_WRITEZF, F_WRITE, F_LOCK, and F_UNLOCK system calls subsequently verify and recompute the checksums when they change the FCB. The F_CLOSE system call also verifies FCB checksums. Note that FCB verification by these system calls can be disabled (see Section 2.12), but Concurrent CP/M's file security is reduced when this is done. If the BDOS detects an FCB checksum error, it does not perform the requested command. Instead, it either returns to the calling process with an error code, or if the system call is F_CLOSE and the BDOS Error mode is in the default state (see Section 2.18), it terminates the calling process with an error message.

Concurrent CP/M uses a system data structure, called the Lock List, to manage file opening and record locking by running processes: Each time a process opens a file or locks a record successfully, the file system allocates an entry in the system Lock List to record the fact. The file system uses the following information to

- prevent a process from deleting, truncating, renaming, or updating the attributes of another process's open file.
- prevent a process from opening a file currently opened by another process, unless both processes open the file in unlocked or Read-Only mode.
- prevent a process from resetting a drive on which another process has an open file.
- prevent a process from reading, writing, or locking a record currently locked by another process. Refer to Section 2.14 for more information on record locking and unlocking.

The file system only verifies whether another process has the FCB-specified file open for the following file-access system calls: F_OPEN, F_MAKE, F_DELETE, F_RENAME, F_ATTRIB, and F_TRUNCATE. For file-access system calls that require an open FCB, the FCB checksum controls whether the calling process can use the FCB. By definition, a valid FCB checksum implies that the file has been successfully opened and an entry for the file resides in the system Lock List.

The most common way a process releases a lock entry for an open file is by closing the file. A close operation is permanent if it causes the removal of the file's open Lock List entry. The file system invalidates the FCB checksum field on permanent close operations to prevent continued open file operations with the FCB.

However, not all close operations are permanent. For example, if a process makes multiple F_OPEN or F_MAKE calls to an open file, a matching number of F_CLOSE calls must be made before the file system permanently closes the file. Of course, if you only open a file once, a single close operation permanently closes the file. In addition, a process can optionally make partial F_CLOSE calls to a file by setting interface attribute F5'. A partial close operation does not affect the open state of a file. In the above example, a partial close operation would not count against an F_OPEN or F_MAKE call. A partial close operation simply updates the directory to reflect the current state of the file.

As a general rule, under Concurrent CP/M a process should close files as soon as it no longer needs them, even if it has not modified them. While a process has a file open, access by other processes to the file is restricted. For example, after a process has opened a file in Locked mode, the file cannot be opened by other processes until the file is closed or the process terminates.

Furthermore, space in the system Lock List is limited. If a process attempts to open a file and no space remains in the system Lock List, or if the process exceeds the open file limit, the BDOS denies the open request and usually terminates the calling process. You can change the way the file system handles this error by making an F_ERRMODE system call. Note that the size of the system Lock List and the process open file limit are GENCCPM parameters.

There are several other situations where the file system removes open file entries from the system Lock List for a process. For example, if a process makes an F_DELETE call for a file it has open in Locked mode, the file system deletes the file and also purges the file's entry from the system Lock List. Deleting an open file is not recommended under Concurrent CP/M but it is supported for files opened in Locked mode to provide compatibility with software written under earlier releases of MP/M™ and CP/M®. The file system does not allow deletion of a file opened in Unlocked or Read-Only mode.

To ensure that the process does not use the open FCB corresponding to the deleted file, the file system subsequently checks all open FCBs for the process. Each open FCB is checked the next time it is used with a file-access system call that requires an open FCB. If a Lock List entry exists for the file, the BDOS allows the operation to proceed; if not, it indicates that the file has been purged and the file system returns an FCB checksum error.

The file system performs this verification of a process's open FCBs whenever it purges an open file entry from the system Lock List. The following list describes these situations:

- A process makes an F_ATTRIB, F_DELETE, F_RENAME, or F_TRUNCATE system call to a file it has open in Locked mode. These operations cannot be performed on a file open in Unlocked or Read-Only mode.
- A process issues a DRV_FREE call for a drive on which it has an open file.
- The BDOS detects a change in media on a drive that has open files. This is a special case because a process cannot control the occurrence of this situation, and because it can impact more than one process. Refer to Section 2.17 for more details on this situation.

Open FCB verification can affect performance because each verification operation requires a directory search operation. In general, you should avoid such situations when creating new programs for Concurrent CP/M.

2.11 Extended File Locking

Extended file locking enables a Concurrent CP/M process to maintain a lock on a file after the file is permanently closed. This facility allows a process to set the attributes, delete, rename, or truncate a file without interference from other processes. In addition, this technique avoids the problems associated with using these system calls on open files (see Section 2.10).

A process can also reopen a file with an extended lock and continue open file processing. To illustrate how extended file locking might be used, a process can close an open file, rename the file, reopen the file under its new name, and continue with file operations without ever losing the file's Lock List item and control over the file.

A process can only specify extended file locking for a file it has opened in Locked mode. To extend a file's lock, set interface attribute F6' when closing the file. The F_CLOSE system call interrogates this attribute only when it is closing a file permanently. Thus, interface attribute F5', signifying a partial close, must be reset when the F_CLOSE call is made. In addition, the close operation must be permanent. If a process has opened a file N times, the F_CLOSE system call ignores the F6' attribute until the file is closed for the Nth time.

Note that the access rules for a file with an extended lock are identical to the rules for a file open in Locked mode. In addition, you cannot extend the lock of a Read-Only file (attribute T1' set), because a Read-Only file cannot be opened in Locked mode.

To maintain an extended file lock through an F_ATTRIB, F_RENAME, or F_TRUNCATE system call, set interface attribute F5' of the referenced FCB when making the call. The BDOS honors this attribute only if the file has been closed with an extended lock. Setting attribute F5' also maintains an extended file lock for the F_DELETE system call, but setting this attribute also changes the nature of the delete operation to an XFCB-only delete. If successful, all four of these system calls delete a file's extended lock item if they are called with attribute F5' reset. However, the extended lock item is not deleted if they return with an error code.

You can make an F_OPEN call to resume record operations on a file with an extended lock. Note that you can also change the open mode when you reopen the file. The following example illustrates the use of extended locks.

1. Open file EXLOCK.TST in Locked mode.
2. Perform read and write operations on the file EXLOCK.TST using the open FCB.
3. Close file EXLOCK.TST with interface attribute F6' set to retain the file's lock item.
4. Use the F_RENAME system call to change the name of the file to EXLOCK.NEW with interface attribute F5' set to retain the file's extended lock item.
5. Reopen the file EXLOCK.NEW in Locked mode.
6. Perform read and write operations on the file EXLOCK.NEW, using the open FCB.
7. Close file EXLOCK.NEW again with interface attribute F6' set to retain the file's lock item.
8. Set the Read-Only attribute and release the file's lock item by making an F_ATTRIB system call with interface attribute F5' reset.

At this point, the file EXLOCK.NEW becomes available for access by another process.

2.12 Compatibility Attributes

Compatibility attributes provide a mechanism to modify some of the Concurrent CP/M file security rules for specific command files. Concurrent CP/M includes this facility because some programs developed under earlier Digital Research operating systems do not run properly under Concurrent CP/M. Most of the problems encountered by these programs occur because they were designed for single-tasking operating systems where file security is not required. For example, a program might close a file and then continue reading and writing to the file. Under CP/M-86, this does not cause a problem. However, under Concurrent CP/M, the file system intercepts open file operations with a deactivated FCB to ensure the integrity of the file system. With compatibility attributes, you have a tool for dealing with these kinds of situations.

You should use compatibility attributes only with existing programs that run properly under CP/M or CP/M-86®. Do not use compatibility attributes with new programs you develop under Concurrent CP/M.

Compatibility attributes are defined as file attributes F1' through F4' of program (CMD) files. You can use the Concurrent CP/M SET utility to set these file attributes from the command line. However, setting a command file's compatibility attributes has no effect unless the GENCCPM COMPATMODE option has been selected during system generation. If this has been done, the P_CLI system call interrogates file attributes F1' through F4' of the command file during program loading and modifies the Concurrent CP/M file security rules for the loaded program.

The Concurrent CP/M BDOS defines the Compatibility Attributes as shown in Table 2-11.

Table 2-11. Compatibility Attribute Definitions

<i>Attribute</i>	<i>Definition</i>
F1'	<p>Modify the rules for Locked mode.</p> <p>When a process running with F1' set opens a file in Locked mode, it can perform read and write operations to the file as normal. However, to other processes on the system, it appears as if the file was opened in Read-Only mode. Thus, another process running with F1' set can open the same file in Locked mode and also perform write operations to the file. In addition, if a process with F1' reset attempts to open the file in Locked or Read-Only mode, the open attempt is allowed but the open mode is forced to Read-Only. Furthermore, write operations are not allowed when the process has F1' reset.</p> <p>This compatibility mode is designed to allow multiple copies of the same program to run concurrently, even though the program might make read and write calls to a common file that it has opened in Locked mode. In addition, this compatibility mode allows other programs not in this compatibility mode to access the file on a Read-Only basis. Note that record locking is not supported for this modified open mode. In addition, to be safe, make all static files such as program and help files Read-Only if you use this compatibility attribute.</p> <p>There is an alternative to using this attribute if a program only makes read calls to the common file. By setting the file's Read-Only attribute, you force the open mode to Read-Only when the file is opened in Locked mode.</p>

Table 2-11. (continued)

<i>Attribute</i>	<i>Definition</i>
F2'	<p>Change F_CLOSE to partial close.</p> <p>Processes running with F2' set only make partial F_CLOSE system calls. This attribute is intended for programs that close a file to update the directory but continue to use the file. A side effect of this attribute is that files opened by a process are not released from the system Lock List until the process terminates. When using this attribute, it might be necessary to set the system Lock List parameters to higher values when you generate a system with GENCCPM.</p>
F3'	<p>Ignore close checksum errors.</p> <p>This attribute changes the way the F_CLOSE system call handles Close Checksum errors. Normally, the file system prints an error message on the console and terminates the calling process. However, if this attribute is set, the F_CLOSE system call ignores the checksum error and performs the close operation. This interface attribute is intended for programs that modify an open FCB before closing a file.</p>
F4'	<p>Disable FCB Checksum verification for read and write operations.</p> <p>Setting this attribute also sets attributes F2' and F3'. This attribute is intended for programs that modify open FCBs during read and write operations. Use this attribute very carefully, and only with software known to work, because it effectively disables Concurrent CP/M's file security.</p>

Use the Concurrent CP/M SET utility to specify the combination of compatibility attributes you want set in the program's command file. For example,

```
A>SET filespec [f1=on]
A>SET filespec [f1=on, f3=on]
A>SET filespec [f4=on]
```

If you have a program that runs under CP/M or CP/M-86 but does not run properly under Concurrent CP/M, use the following guidelines to select the proper compatibility attributes for the program.

- If the program ends with the “File Currently Opened” message when multiple copies of the program are run, set compatibility attribute F1', or place all common static files under User 0 with the SYS and Read-Only attributes set.
- If the program terminates with the message “Close Checksum Error”, set compatibility attribute F3'.
- If the program terminates with an I/O error, try running the program with attribute F2' set. If the problem persists, then try attribute F4'. Use attribute F4' only as a last resort.

2.13 Multisector I/O

The BDOS file system provides the capability to read or write multiple 128-byte records in a single BDOS system call. This multisector facility can be visualized as a BDOS burst mode, enabling a process to complete multiple I/O operations without interference from other running processes. In addition, the BDOS file system bypasses, when possible, all intermediate record buffering during multisector I/O operations. Data is transferred directly between the calling process's memory and the drive. The BDOS also informs the XIOS when it is reading or writing multiple physical records on a drive. The XIOS can use this information to further optimize the I/O operation resulting in even better performance. As a result, the use of this facility in an application program can improve its performance and also enhance overall system throughput, particularly when performing sequential I/O.

The number of records that can be transferred with multisector I/O ranges from 1 to 128. This value, called the BDOS Multisector Count, can be set by the F_MULTISEC system call. The P_CLI system call sets the Multisector Count to 1 when it initiates a transient program for execution. Note that the greatest potential performance increases are obtained when the Multisector Count is set to 128. Of course, this requires a 16K buffer. The Concurrent CP/M PIP utility performs its sequential I/O with a Multisector Count of 128.

The Multisector Count determines the number of operations to be performed by the following BDOS system calls:

- F_READ and F_WRITE system calls
- F_READRAND, F_WRITERAND, and F_WRITEZF
- F_LOCK and F_UNLOCK

If the Multisector Count is N , calling one of the above system calls is equivalent to making N system calls. With the exception of disk I/O errors encountered by the XIOS, if an error interrupts a multisector read or write operation, the file system returns the number of 128-byte records successfully transferred in register AH. Section 2.14 describes how the Multisector Count affects the F_LOCK and F_UNLOCK system calls.

2.14 Concurrent File Access

Concurrent CP/M supports two open modes, Read-Only and Unlocked, which allow concurrently running processes to access common files for record operations. The Read-Only open mode allows multiple processes to read from a common file, but processes cannot write to a file open in this mode. Thus, files remain static when they are opened in Read-Only mode. The Unlocked open mode is more complex because it allows multiple processes to read and write records to a common file. As a result, Unlocked mode has some important differences from the other open modes.

When a process opens a file in Unlocked mode, the file system returns a 2-byte field called the File ID in the R0 and R1 bytes of the FCB. The File ID is a required parameter of Concurrent CP/M's record locking system calls, F_LOCK and F_UNLOCK, which are only supported for files open in Unlocked mode. Note that these system calls return a successful error code if they are called for files opened in Locked mode. However, they perform no action in this case, because, by definition, the calling process has the entire file locked.

The F_LOCK and F_UNLOCK system calls allow a process to establish and release temporary ownership to particular records within a file. You must set the FCB Random Record field and place the File ID in the first two bytes of the current DMA buffer before making these calls. The file system locks and unlocks records in units of 128 bytes, which is the standard Concurrent CP/M record size. The number of records locked or unlocked is controlled by the BDOS Multisector Count, which can range from 1 to 128 (see Section 2.13). In order to simplify the discussion of record locking and unlocking, the following paragraphs assume the Multisector Count is one. However, as discussed later in this section, the more general case of multiple record locking and unlocking is a simple extension of the single record case.

The F_LOCK system call supports two types of lock operations: exclusive locks and shared locks. Interface attribute F5' specifies the type of lock. F5' = 0 requests an exclusive lock; F5' = 1 requests a shared lock. If a process locks a record with an exclusive lock, other processes cannot read, write, or lock the record. The locking process, however, can access the record with no restrictions. You should use this type of lock when exclusive control over a record is required.

If a process locks a record with a shared lock, other processes cannot write to the record or make an exclusive lock of the record. However, other processes are allowed to read the record and make their own shared locks on the record. No process, including the locking process, can write to a record with a shared lock. Shared locks are useful when you want to ensure that a record does not change, but you want to allow other processes to read the record.

The F_LOCK system call also lets you change the lock of a record if there is no conflict. For example, you can convert an exclusive lock into a shared lock with no restrictions. On the other hand, a process cannot convert a record's shared lock to an exclusive lock if another process has a shared lock on the record.

The F_LOCK system call has another option, specified by interface attribute F6', which controls whether a record must exist in order to be locked. If you make an F_LOCK system call with F6' = 0, the file system returns an error code if the specified record does not exist within the file. Setting F6' to 1 requests a logical lock operation. Logical lock operations are only limited by the maximum Concurrent CP/M file size of 32 megabytes, which corresponds to a maximum Random Record Number of 262,143. You can use logical locks to control extending a shared file.

The F_UNLOCK system call is similar to the F_LOCK call except that it removes locks instead of creating them. There are few restrictions on unlock operations. Of course a process can only remove locks that it has made. The F_UNLOCK system call has one option, controlled by interface attribute F5'. If F5' is set to one, the F_UNLOCK system call removes all locks for the file made by the calling process. Otherwise, it removes the locks specified by the Random Record field and the BDOS Multisector Count. Note that the F_CLOSE system call also removes all locks for a file on permanent close operations.

If the BDOS Multisector Count is greater than one, the F_LOCK and F_UNLOCK system calls perform multiple record locking or unlocking. In general, multiple record locking and unlocking can be viewed as a sequence of N independent operations, where N equals the Multisector Count. However, if an error occurs on any record within the sequence, no locking or unlocking is performed. For example, both F_LOCK and F_UNLOCK perform no action and return an error code if the sum of the FCB Random Record Number and the BDOS Multisector Count is greater than 262,144. As another example, the F_LOCK system call also returns an error code if another process has an exclusive lock on any record within the sequence.

When a process makes an F_LOCK system call, the file system allocates a new entry in the system Lock List to record the lock operation and associate it with the calling process. A corresponding F_UNLOCK system call removes the locked entry from the list. While the lock entry exists in the system Lock List, the file system enforces the restrictions implied by the lock item.

Because each lock item includes a record count field, a multiple lock operation normally results in the creation of a single new entry. However, if the file system must split an existing lock entry to satisfy the lock operation, an additional entry is required. Similarly, an unlock operation can require the creation of a new entry if a split is needed. Thus, in the worst case, a lock operation can require two new lock entries and an unlock operation can require one. Note that lock item splitting can be avoided by locking and unlocking records in consistent units.

These considerations are important because the Lock List is a limited resource under Concurrent CP/M. The file system performs no action and returns an error code if insufficient available entries exist in the system Lock List to satisfy the lock or unlock request. In addition, the number of lock items a single process is allowed to consume is a GENCCPM parameter established at SYSGEN time. The file system also returns an error code if this limit is exceeded.

The file system performs several special operations for read and write system calls to a file open in Unlocked mode. These operations are required because the file system maintains the current state of an open file in the calling process's FCB. When multiple processes have the same file open, FCBs for the same file exist in each process's memory. To ensure that all processes have current information, the file system updates the directory immediately when an FCB for an unlocked file is changed. In addition, the file system verifies error situations such as end-of-file, or reading unwritten data with the directory before returning an error. As a result, read and write operations are less efficient for files open in Unlocked mode when compared to equivalent operations for files opened in Locked mode.

2.15 File Byte Counts

Although the logical record size of Concurrent CP/M is restricted to 128 bytes, the file system does provide a mechanism to store and retrieve a byte count for a file. This facility can identify the last byte of the last record of a file. The F_SIZE system call returns the Random Record Number, + 1, of the last record of a file.

The F_ATTRIB system call can set a file's byte count. This is an option controlled by interface attribute F6'. Conversely, the F_OPEN system call can return a file's byte count to the CR field of the FCB. The F_SFIRST and F_SNEXT system calls also return a file's byte count. These system calls return the byte count in the CS field of the FCB returned in the current DMA buffer.

Note that the file system does not access or update the byte count value in BDOS read or write system calls. However, the F_MAKE system call does set the byte count value to zero when it creates a file in the directory.

2.16 Record Blocking and Deblocking

Under Concurrent CP/M, the logical record size for disk I/O is 128 bytes. This is the basic unit of data transfer between the operating system and running processes. However, on disk, the record size is not restricted to 128 bytes. These records, called physical records, can range from 128 bytes to 4K bytes in size. Record blocking and deblocking is required on systems that support drives with physical record sizes larger than 128 bytes.

The process of building up physical records from 128-byte logical records is called record blocking. This process is required in write operations. The reverse process of breaking up physical records into their component 128-byte logical records is called record deblocking. This process is required in read operations. Under Concurrent CP/M, record blocking and deblocking is normally performed by the BDOS.

Record deblocking implies a read-ahead operation. For example, if a process reads a logical record that resides at the beginning of a physical record, the entire physical record is read into an internal buffer. Subsequent BDOS read calls for the remaining logical records access the buffer instead of the disk. Conversely, record blocking results in the postponement of physical write operations but only for data write operations. For example, if a transient program makes a BDOS write call, the logical record is placed in a buffer equal in size to the physical record size. The write operation on the physical record buffer is postponed until the buffer is needed in another I/O operation. Note that under Concurrent CP/M, directory write operations are never postponed.

Postponing physical record write operations has implications for some application programs. For programs that involve file updating, it is often critical to guarantee that the state of the file on disk parallels the state of the file in memory after an update operation. This is only an issue on drives where physical write operations are postponed because of record blocking and deblocking. If the system should crash while a physical buffer is pending, data would be lost. To prevent this loss of data, the `F_FLUSH` system call can be called to force the write of any pending physical buffers associated with the calling process.

Note: The file system discards all pending physical data buffers when a process terminates. However, the file system automatically makes an `F_FLUSH` call in the `F_CLOSE` system call. Thus, it is sufficient to make an `F_CLOSE` system call to ensure that all pending physical buffers for that file are written to the disk.

2.17 Reset, Access, and Free Drive

The BDOS system calls `DRV_ALLRESET`, `DRV_RESET`, `DRV_ACCESS`, and `DRV_FREE` allow a process to control when to reinitialize a drive directory for file operations. This process of initializing a drive's directory is called logging-in the drive.

When you start Concurrent CP/M, all drives are initialized to the reset state. Subsequently, as processes reference drives, the file system automatically logs them in. Once logged-in, a drive remains in the logged-in state until it is reset by the `DRV_ALLRESET` or `DRV_RESET` system calls or a media change is detected on the drive. If the drive is reset, the file system automatically logs in the drive again the next time a process references it. The file system logs in a drive immediately when it detects a media change on the drive.

Note that the `DRV_ALLRESET` and `DRV_RESET` system calls have similar effects except that the `DRV_ALLRESET` system call affects all drives on the system. You can specify the combination of drives to reset with the `DRV_RESET` system call.

Logging-in a drive consists of several steps. The most important step is the initialization of the drive's allocation vector. The allocation vector records the allocation and deallocation of data blocks to files, as files are created, extended, deleted and truncated. Another function performed during drive log-in is the initialization of the directory checksum vector. The file system uses the checksum vector to detect media changes on a drive. Note that permanent drives, which do not support media changes, usually do not have checksum vectors.

Under Concurrent CP/M, the DRV_RESET operation is conditional. The file system cannot reset a drive for a process if another process has an open file on the drive. However, the exact action taken by a DRV_RESET operation depends on whether the drive to be reset is permanent or removable.

Concurrent CP/M determines whether a drive is permanent or removable by interrogating a bit in the drive's Disk Parameter Block (DPB) in the XIOS. A high-order bit of 1 in the DPB Checksum Vector Size field designates the drive as permanent. A drive's Removable or Nonremovable designation is critical to the reset operation described below.

The BDOS first determines whether there are any files currently open on the drive to be reset. If there are none, the reset takes place. If there are open files, the action taken by the reset operation depends on whether the drive is removable and whether the drive is Read-Only or Read-Write. Note that only the DRV_SETRO system call can set a drive to Read-Only. Following log-in, a drive is always Read-Write.

If the drive is a permanent drive and if the drive is not Read-Only, the reset operation is not performed, but a successful result is returned to the calling process.

However, if the drive is removable or set to Read-Only, the file system determines whether other processes have open files on the drive. If they do, then it denies DRV_RESET operation and returns an error code to the calling process.

If all the open files on a removable drive belong to the calling process, the process is said to own the drive. In this case, the file system performs a qualified reset on the drive and returns a successful result. This means that the next time a process accesses this drive, the BDOS performs the log-in operation only if it detects a media change on the drive. The logic flow of the drive reset operation is shown in Figure 2-6.




Figure 2-6. Disk System Reset

If the BDOS detects a media change on a drive after a qualified reset, it purges all open files on the drive from the system Lock List and subsequently verifies all open FCBs in file operations for the owning process (refer to Section 2.10 for details of FCB verification).

In all other cases where the BDOS detects a media change on a drive, the file system purges all open files on the drive from the system Lock List, and flags all processes owning a purged file for automatic open FCB verification.

Note: If a process references a purged file with a BDOS command that requires an open FCB, the file system returns to the process with an FCB checksum error.

The primary purpose of the drive reset functions is to prepare for a media change on a drive. Because a drive reset operation is conditional, it allows a process to test whether it is safe to change disks. Thus, a process should make a successful drive reset call before prompting the user to change disks. In addition, you should close all your open files on the drive, particularly files you have written to, before prompting the user to change disks. Otherwise, you might lose data.

The `DRV_ACCESS` and `DRV_FREE` system calls perform special actions under Concurrent CP/M. The `DRV_ACCESS` system call inserts a dummy open file item into the system Lock List for each specified drive. While that item exists in the system Lock List, no other process can reset the drive. The `DRV_FREE` system call purges the Lock List of all items, including open file items, belonging to the calling process on the specified drives. Any subsequent reference to those files by a BDOS system call requiring an open FCB results in an FCB checksum error return.

The `DRV_FREE` system call has two important side effects. First of all, any pending blocking/deblocking buffers on a specified drive that belong to the calling process are discarded. Secondly, any data blocks that have been allocated to files that have not been closed are lost. Be sure to close your files before making this system call.

The `DRV_SETRO` system call is also conditional under Concurrent CP/M. The file system does not allow a process to set a drive to Read-Only if another process has an open file on the drive. This applies to both removable and permanent drives.

A process can prevent other processes from resetting a Read-Only drive by opening a file on the drive or by issuing a `DRV_ACCESS` call for the drive and then making a `DRV_SETRO` system call. Executing `DRV_SETRO` before the `F_OPEN` or `DRV_ACCESS` call leaves a window in which another process could set the drive back to Read-Write. While the open file or dummy item belonging to the process resides in the system Lock List, no other process can reset the drive to take it out of Read-Only status.

2.18 BDOS Error Handling

The Concurrent CP/M file system has an extensive error handling capability. When an error is detected, the BDOS responds in one of three ways:

1. It can return to the calling process with return codes in the AX register identifying the error.
2. It can display an error message on the console and terminate the process.
3. It can display an error message on the console and return an error code to the calling process, as in method 1.

The file system handles the majority of errors it detects by method 1. Two examples of this kind of error are the "file not found" error for the F_OPEN system call and the "reading unwritten data" error for the F_READ call. More serious errors, such as disk I/O errors, are normally handled by method 2. Errors in this category, called physical and extended errors, can also be reported by methods 1 and 3 under program control.

The BDOS Error mode, which has three states, determines how the file system handles physical and extended errors. In the default state, the BDOS displays the error message and terminates the calling process (method 2). In Return Error mode, the BDOS returns control to the calling process with the error identified in the AX register (method 1). In Return and Display Error mode, the BDOS returns control to the calling process with the error identified in the AX register and also displays the error message at the console (method 3).

While both return modes protect a process from termination because of a physical or extended error, the Return and Display mode also allows the calling process to take advantage of the built-in error reporting of the file system. Physical and extended errors are displayed on the console in the following format:

```
CP/M Error on d: error message  
BDOS Function = nn File = filename.typ
```

where d is the name of the drive selected when the error condition occurs; error message identifies the error; nn is the BDOS function number, and filename.typ identifies the file specified by the BDOS function. If the BDOS function did not involve an FCB, the file information is omitted.

Tables 2-12 and 2-13 detail BDOS physical and extended error messages.

Table 2-12. BDOS Physical Errors

<i>Message</i>	<i>Meaning</i>
Disk I/O	The "Disk I/O" error results from an error condition returned to the BDOS from the XIOS module. The file system makes XIOS read and write calls to execute BDOS file-access system calls. If the XIOS read or write routine detects an error, it returns an error code to the BDOS, causing this error message.
Invalid Drive	The "Invalid Drive" error also results from an error condition returned to the BDOS from the XIOS module. The BDOS makes an XIOS Select Disk call before accessing a drive to perform a requested BDOS function. If the XIOS does not support the selected disk, it returns an error code resulting in this error.
Read/Only File	The BDOS returns the "Read/Only File" error message when a process attempts to write to a file with the R/O attribute set.
Read/Only Disk	The BDOS returns the "Read/Only Disk error" message when a process makes a write operation to a disk that is in Read-Only status. A drive can be placed in Read-Only status explicitly with the DRV_SETRO system call.

Table 2-13. BDOS Extended Errors

<i>Message</i>	<i>Meaning</i>
File Opened in Read/Only Mode	<p>The BDOS returns the "File Opened in Read/Only Mode" error message when a process attempts to write to a file opened in Read-Only mode. A process can open a file in Read-Only mode explicitly by setting FCB interface attribute F6'. In addition, if a process opens a file in Locked mode, the file system automatically forces the open mode to Read-Only mode when:</p> <ul style="list-style-type: none"> ■ the process opens a file with the Read-Only attribute set. ■ the current user number is not zero and the process opens a user zero file with the SYS attribute set. <p>The BDOS also returns this error if a process attempts to write to a file that is password-protected in Write mode, and it did not supply the correct password when it opened the file.</p>
File Currently Open	<p>The BDOS returns the "File Currently Open" error message when a process attempts to delete, rename, or modify the attributes of a file opened by another process. The BDOS also returns this error when a process attempts to open a file in a mode incompatible with the mode in which the file was previously opened by another process or by the calling process.</p>
Close Checksum Error	<p>The BDOS returns the "Close Checksum Error" message when the BDOS detects a checksum error in the FCB passed to the file system with an F_CLOSE call.</p>
Password Error	<p>The BDOS returns the "Password Error" message when passwords are required and the file password is not supplied or is incorrect.</p>

Table 2-13. (continued)

<i>Message</i>	<i>Meaning</i>
File Already Exists	The BDOS returns the "File Already Exists" error message for the F_MAKE and F_RENAME system calls when the BDOS detects a conflict on filename and filetype.
Illegal ? in FCB	The BDOS returns the "Illegal ? in FCB" error message when the BDOS detects a ? character in the filename or filetype of the passed FCB for the F_ATTRIB, F_OPEN, F_RENAME, F_TIMEDATE, F_WRITEFCB, F_TRUNCATE, and F_MAKE system calls.
Open File Limit Exceeded	The BDOS returns the "Open File Limit Exceeded" error message when a process exceeds the process file lock limit specified by GENCCPM. The F_OPEN, F_MAKE, and DRV_ACCESS system calls can return this error.
No Room in System Lock List	The BDOS returns the "No Room in System Lock List" error message when no room for new entries exists within the system Lock List. The F_OPEN, F_MAKE, and DRV_ACCESS system calls can return this error.

The following paragraphs describe the error return code conventions of the file system calls. Most file system calls fall into three categories in regard to return codes; they return an error code, a directory code, or an error flag. The error conventions let programs written for CP/M-86 run without modification.

The following BDOS system calls return a logical error in register AL:

- F_LOCK
- F_READ
- F_READRAND
- F_UNLOCK
- F_WRITE
- F_WRITERAND
- F_WRITEZF

Table 2-14 lists error code definitions for register AL.

Table 2-14. BDOS Error Codes

<i>Code</i>	<i>Definition</i>
00H:	Function successful
01H:	Reading unwritten data
02H:	No available directory space (Write Sequential)
03H:	No available data block
04H:	Cannot close current extent
05H:	Seek to unwritten extent
06H:	No available directory space
07H:	Random record number out of range
* 08H:	Record locked by another process (restricted to files opened in Unlocked mode)
09H:	Invalid FCB (previous BDOS F_CLOSE system call returned an error code and invalidated the FCB)
0AH:	FCB checksum error
* 0BH:	Unlocked file unallocated block verify error
** 0CH:	Process record lock limit exceeded
** 0DH:	Invalid File ID
** 0EH:	No room in System Lock List
0FFH:	Physical error : refer to register AH
* - returned only for files opened in Unlocked mode	
** - returned only by the F_LOCK and F_UNLOCK system calls for files opened in Unlocked mode	

For BDOS read and write system calls, the file system also sets register AH when the returned error code is a value other than zero or 0FFH. In this case, register AH contains the number of 128-byte records successfully read or written before the error was encountered. Note that register AH can only contain a nonzero value if the calling process has set the BDOS Multisector Count to a value other than one; otherwise register AH is always set to zero. On successful system calls (Error Code = 0), register AH is also set to zero. If the Error Code is 0FFH, register AH contains a physical error code (see Table 2-15).

The following BDOS system calls return a directory code in register AL:

DRV_SETLABEL
F_ATTRIB
F_CLOSE
F_DELETE
F_MAKE
F_OPEN
F_RENAME
F_SIZE
F_SFIRST
F_SNEXT
F_TIMEDATE
F_TRUNCATE
F_WRITEXFCB

The directory code definitions for register AL follow.

00H - 03H : successful function
0FFH : unsuccessful function

With the exception of the F_SFIRST and F_SNEXT system calls, all functions in this category return with the directory code set to zero upon a successful return. However, for these two system calls, a successful directory code identifies the relative starting position of the directory entry in the calling process's current DMA buffer.

If a process uses the F_ERRMODE system call to place the BDOS in Return Error mode, the following system calls return an error flag in register AL on physical errors:

DRV_GETLABEL
 DRV_ACCESS
 DRV_SET
 DRV_SPACE
 DRV_FLUSH

The error flag definition for register AL follows.

00H : successful function
 0FFH : physical error : refer to register AH

The BDOS returns nonzero values in register AH to identify a physical or extended error if the BDOS Error mode is in one of the return modes. Except for system calls that return a Directory Code, register AL equal to 0FFH indicates that register AH identifies the physical or extended error. For functions that return a Directory Code, if register AL equals 255, and register AH is not equal to zero, register AH identifies the physical or extended error. Table 2-15 shows the physical and extended error codes returned in register AH.

Table 2-15. BDOS Physical and Extended Errors

<i>Code</i>	<i>Explanation</i>
01H	Disk I/O Error : permanent error
02H	Read/Only Disk
03H	Read/Only File, File Opened in Read/Only Mode, or File Password Protected in Write Mode and Correct Password Not Specified
04H	Invalid Drive : drive select error
05H	File Currently Open in an incompatible mode
06H	Close Checksum Error
07H	Password Error
08H	File Already Exists
09H	Illegal ? in FCB
0AH	Open File Limit Exceeded
0BH	No Room in System Lock List

The following two system calls represent a special case because they return an address in register AX.

DRV_ALLOCVEC
DRV_DBP

When the calling process is in one of the BDOS return error modes and the BDOS detects a physical error for these system calls, it returns to the calling process with registers AX and BX set to 0FFFFH. Otherwise, they return no error code.

Under Concurrent CP/M, the following system calls also represent a special case.

DRV_ALLRESET
DRV_RESET
DRV_SETRO

These system calls return to the calling process with registers AL and BL set to 0FFH if another process has an open file or has made a DRV_ACCESS call that prevents the reset or write protect operation. If the calling process is not in Return Error mode, these system calls also display an error message identifying the process that prevented the requested operation.

End of Section 2

Section 3

Transient Commands

3.1 Transient Program Load and Exit

A transient program is a file of type CMD that is loaded from disk and resides in memory only during its operation. A resident system program is a file of type RSP that is included in Concurrent CP/M during GENCCPM. Section 4 describes the three system memory models that determine the initial values of segment registers in transient processes.

You can initiate a transient process by entering a command at a system console. The console's TMP (Terminal Message Processor) then calls the Command Line Interpreter system call (refer to the P_CLI system call), and passes to it the command line entered by the user. If the command is not an RSP, then the P_CLI system call locates and then loads the proper CMD file. P_CLI then calls the F_PARSE system call to parse up to two filenames following the command, and place the properly formatted FCBs at locations 005CH and 006CH in the Base Page of the initial Data Segment.

The P_CLI system call initializes memory, the Process Descriptor, and the User Data Area (UDA), and allocates a 96-byte stack area, independent of the program, to contain the process's initial stack. If 8087 processing is required (see Section 3.1.2) P_CLI allocates an additional 96 bytes for the UDA. Concurrent CP/M divides the DMA address into the DMA segment address and the DMA offset. P_CLI initializes the default DMA segment to the value of the initial data segment, and the default DMA offset to 0080H.

The P_CLI system call creates the new process with a P_CREATE system call and sets the initial stack so that the process can execute a Far Return instruction to terminate. A process also ends when it calls DRV_ALLRESET or P_TERM.

You can also terminate a process by typing a single CTRL-C during console input. See C_MODE for details of enabling/disabling CTRL-C. CTRL-C, when typed at the prompt, forces a DRV_RESET call for each logged-in drive. This operation only affects removable media drives.

Note: Additional UDA space is allocated for 8087 processing only if the process is initialized by the P_CLI or P_LOAD system call. Other processes (such as RSPs) that require 8087 processing and do not use P_CLI or P_LOAD must allocate this additional UDA space themselves.

3.1.1 Shared Code

Concurrent CP/M allows processes to share program code. This capability of sharing program code avoids unnecessary program loading of a code segment already in memory and conserves memory space since multiple copies of the same program code do not have to occupy different memory space. During program load of a "sharable" program code, the system allocates the code group separately from the rest of the program. This code group is maintained in memory even after the program has terminated. Subsequent loading of the same program does not load the code group, but uses the existing one instead. Obviously, programs written with separate code and data can take advantage of this feature.

The system maintains a shared code group in memory until a memory request or a reset drive forces its release. The system maintains shared code groups in memory in Least Recently Used (LRU) order on the Shared Code List. If a memory request is made that cannot be satisfied, the list is drained, one at a time, until the memory request is satisfied, or the Shared Code List is emptied. If a drive is reset, the system purges all code groups from the Shared Code List loaded from that drive.

A shared code program is flagged by the value 09H in the G_Type field of the Code Group Descriptor in the CMD file header (see Section 3.2). The user may set this field by using the CHSET utility (see *Concurrent CP/M User's Guide*). Note that programs using the 8080 memory model cannot be set to shared code.

3.1.2 8087 Support

Concurrent CP/M provides optional 8087 support for systems that use the 8087 processor. This support is indicated by the Program Flag, byte 127 (07FH), of the CMD file header. Setting bit 6 (bit 0 is least significant bit) of the Program Flag indicates optional 8087 support, which means that if the 8087 is present, the program uses it; otherwise, the program will emulate it. If bit 5 of the Program Flag is set, it indicates that the 8087 must be present in order for the program to run. If no 8087 is present and bit 5 of the Program Flag is set, the system returns an error when it tries to load the program. The CHSET utility can be used to set the program's header record for optional or required 8087 support.

If you use the P_CLI or P_LOAD system call to initiate and execute a process, the system allocates an extra 96 bytes to the UDA for 8087 support. If you require 8087 support and do not use the P_CLI or P_LOAD system call, you must specifically allocate this additional 96 bytes to the UDA, turn on the 8087 flag in the PD, and initialize the CW and SW fields in the 8087 UDA extension (see description of these fields in Section 6 under the P_CREATE system call).

3.1.3 8087 Exception Handling

Although the system provides its own 8087 exception handling routine, the user might want to write his own 8087 exception handler. Appendix E includes instructions and information required by the user to write his own 8087 exception handler, with a sample listing of an 8087 exception handler routine.

3.2 Command File Format

A CMD file consists of a 128-byte header record followed immediately by the memory image. The command file header record is composed of 8 group descriptors (GDs), each 9 bytes long. Each group descriptor describes a portion of the program to be loaded. The format of the header record is shown in Figure 3-1.




Figure 3-1. CMD File Header Format

In Figure 3-1, GD 1 through GD 8 represent group descriptors. Each group descriptor corresponds to an independently loaded program unit and has the format shown in Figure 3-2.




Figure 3-2. Group Descriptor Format

G_Type determines the group descriptor type. The valid group descriptors have a G_Type in the range 1 through 8, as shown in Table 3-1. All other values are reserved for system use. For a given CMD file header only a Code Group and one of any other type can be included.

If a program uses either the Small or Compact Model, the code group is typically pure; that is, it is not modified during program execution.

Table 3-1. Group Descriptors

<i>G_Type</i>	<i>Group Type</i>
01H	Code Group (non-shared)
02H	Data Group
03H	Extra Group
04H	Stack Group
05H	Auxiliary Group #1
06H	Auxiliary Group #2
07H	Auxiliary Group #3
08H	Auxiliary Group #4
09H	Code Group (shared)

All remaining values in the group descriptor are given in increments of 16-byte paragraph units with an assumed low-order 0 nibble to complete the 20-bit address.

Table 3-2. Group Descriptor Fields

<i>Field</i>	<i>Description</i>
G_Length	Gives the number of paragraphs in the group. Given a G_length of 080H, for example, the size of the group is 0800H (2048 decimal) bytes.
A_Base	Defines the base paragraph address for a nonrelocatable group.
G_Min/G_Max	Define the minimum and maximum size of the memory area to allocate to the group.

The memory model described by a header record is implicitly determined by the group descriptors (refer to Section 4.1). The 8080 Model is assumed when only a code group is present, because no independent data group is named. The Small Model is assumed when both a code and data group are present but no additional group descriptors occur. Otherwise, the Compact Model is assumed when the CMD file is loaded.

3.3 Base Page Initialization

The Concurrent CP/M Base Page contains default values and locations initialized by the P_CLI and P_LOAD system calls and used by the transient process.

The Base Page occupies the regions from offset 0000H through 00FFH relative to the initial data segment, and contains the values shown in Figure 3-3.

	L	M	H	L	H	
0	1	2	3	4	5	6
0	CODE LENGTH			CODE BASE		M80
6	DATA LENGTH			DATA BASE		RESERVED
C	EXTRA LENGTH			EXTRA BASE		RESERVED
12	STACK LENGTH			STACK BASE		RESERVED
18	AUX 1			AUX 1		RESERVED
1E	AUX 2			AUX 2		RESERVED
24	AUX 3			AUX 3		RESERVED
2A	AUX 4			AUX 4		RESERVED
30	BYTES 030H THROUGH 04FH ARE NOT CURRENTLY USED AND ARE RESERVED FOR FUTURE USE BY DIGITAL RESEARCH . . .					
50	DRIVE	PASSWORD 1 ADDR		P1 LEN	PASSWORD 2 ADDR	
56	P2 LEN	RESERVED FOR FUTURE USE				
5C	DEFAULT FILE NAME1 . .					
6C	DEFAULT FILE NAME2 . .					
7C	CR	RANDOM RECORD NUMBER (OPT)				
80	DEFAULT 128-BYTE DMA BUFFER					

Figure 3-3. Concurrent CP/M Base Page Values

The fields in the Base Page are defined as follows:

- The M80 byte is a flag indicating whether the 8080 Memory Model was used during load. The values of the flag are defined as:

1 = 8080 Model

0 = not 8080 Model

If the 8080 Model is used, the code length never exceeds 0FFFFH.

- The bytes marked Aux 1 through Aux 4 correspond to a set of four optional independent groups that might be required for programs that execute using the Compact Memory Model. The initial values for these descriptors are derived from the header record in the memory image file.
 - Length is stored using the Intel convention: low, middle, and high bytes.
 - Base refers to the paragraph address of the beginning of the segment.
 - The drive byte identifies the drive from which the transient program was read. 0 designates the default drive, while a value of 1 through 16 identifies drives A through P.
 - Password 1 Addr (bytes 0051H-0052H) contains the address of the password field of the first command tail operand in the default DMA buffer at 0080H. The P_CLI system call sets this field to 0 if no password is specified.
 - P1 Len (byte 0053H) contains the length of the password field for the first command tail operand. The P_CLI system call sets this to 0 if no password is specified.
 - Password 2 Addr (bytes 0054H-0055H) contains the address of the password field of the second command tail operand in the default DMA buffer at 0080H. The P_CLI system call sets this field to 0 if no password is specified.
 - P2 Len (byte 0056H) contains the length of the password field for the second command tail operand. The P_CLI system call sets this field to 0 if no password is specified.
 - File Name1 (bytes 005CH-0067H) is initialized by the P_CLI system call for a transient program from the first command tail operand of the command line.
 - File Name2 (bytes 006CH-0077H) is initialized by the P_CLI system call for a transient program from the second command tail operand of the command line.
- Note:** File Name1 can be used as part of a File Control Block (FCB) beginning at 05CH. To preserve File Name2, copy it to another location before using the FCB in file I/O system calls.
- The CR field (byte 007CH) contains the current record position used in sequential file operations with the FCB at 05CH.

- The optional Random Record Number (bytes 007DH-007FH) is an extension of the FCB at 05CH, used in random record processing.
- The Default DMA buffer (bytes 0080H-00FFH) contains the command tail when the P_CLI system call loads a transient program.

3.4 Parent/Child Relationships

Under Concurrent CP/M when one process creates another process, there is a parent/child relationship between them. The child process inherits most of the default values of the parent process. This includes the default disk, user number, console, list device, and password. The child process also inherits interrupt vectors 0, 1, 3, 4, 224, and 225, which the parent process initialized.

3.5 Direct Video Mapping

Processes which bypass Concurrent CP/M Character I/O system calls and use a video map or screen buffer directly cannot be monitored by the system and continue to display characters on the screen even when running in the background. Consequently, any screen displayed by the program in the foreground console is interspersed with characters displayed by the program in the background using direct video map I/O. To avoid the screen problems created by using direct video I/O, set bit 3 of the Program Flag to indicate to the system that the process is to be put in suspend mode whenever it is running in the background and may continue running only when it is switched to the foreground. The CHSET utility (see the *Concurrent CP/M User's Guide*) can be used to set bit 3 of the Program Flag.

Note that bypassing the system Character I/O system calls negates the concurrency of a process, since the system suspends it from running (if bit 3 of Program Flag is set) unless it is running in the foreground.

End of Section 3

Section 4

Command File Generation

4.1 Transient Execution Models

When the program is loaded, the initial values of the segment registers, the instruction pointer, and the stack pointer are determined by the specific type of memory model used by the transient process, indicated in the CMD file header record.

There are three memory models, the 8080 model, the Small Model, and the Compact Model, summarized in Table 4-1.

Table 4-1. Concurrent CP/M Memory Models

<i>Model</i>	<i>Group Relationships</i>
8080 Model	Code and Data Groups Overlap
Small Model	Independent Code and Data Groups
Compact Model	Three or More Independent Groups

The 8080 Model supports programs that are directly translated from an 8080 environment where code and data are intermixed. The 8080 Model consists of one group that contains all the code, data, and stack areas. Segment registers are initialized to the starting address of the region containing this group. The segment registers can, however, be managed by the application program during execution so that multiple segments in the code group can be addressed.

The Small Model is similar to that defined by Intel, where the program consists of an independent code group and a data group. The code and data groups often consist of, but are not restricted to, single 64K byte segments.

The Compact Model occurs when any of the extra, stack, or auxiliary groups are present in program. Each group can consist of one or more segments, but if any group exceeds one segment in size, or if auxiliary groups are present, then the application program must manage its own segment registers during execution in order to address all code and data areas.

These three models differ primarily in how the operating system initializes the segment registers when it loads a transient process. The P_LOAD system call determines the memory model used by a transient program by examining the program group usage, as described in the following sections.

For all models, the system initializes an internal 96-byte stack area. The first two words of this stack are reserved for the double word return for termination by a RETF (Far Return) instruction. The initial program stack for all models is shown in Figure 4-1 below.




Figure 4-1. Initial Program Stack

The transient program can terminate by using the P_TERMCPM or P_TERM system call or by executing a RETF (Far Return) instruction when the SS and SP still point to the initial program stack.

4.1.1 The 8080 Memory Model

The 8080 Model is assumed when the transient program contains only a code group. In this case, the Command Line Interpreter (P_CLI) system call initializes the CS, DS, and ES registers to the beginning of the code group and sets the SS and SP registers to a 96-byte initial stack area that it allocates.

Note: The P_CLI system call initializes the stack so that if the process executes a Far Return instruction, it terminates. This system call sets the Instruction Pointer (IP) Register to 100H, thus allowing Base Page values at the beginning of the code group. Following program load, the 8080 Model appears as shown in Figure 4-2.




Figure 4-2. Concurrent CP/M 8080 Memory Model

The intermixed code and data areas are indistinguishable. The Base Page values are described in Section 3.3. The following ASM-86 example shows how to code an 8080 Model transient assembly language program.

```

cseg
org 100h
.
. (code)
endcs equ $
dseg
org offset endcs
.
. (data)
end
    
```

4.1.2 The Small Memory Model

The Small Model is assumed when the transient program contains both a code and data group. (In ASM-86, all code is generated following a CSEG directive. Data is defined following a DSEG directive, with the origin of the Data Segment independent of the Code Segment.) In this model, the P_CLI system call sets the CS register to the beginning of the code group, the IP to 0000H, the DS and ES registers to the beginning of the data group, and the SS and SP registers to a 96-byte initial stack area that it initializes. Following program load, the Small Model appears as shown in Figure 4-3.




Figure 4-3. Concurrent CP/M Small Memory Model

The machine code begins at CS + 0000H, the Base Page values begin at DS + 0000H, and the data area starts at DS + 0100H. The following ASM-86 example shows how to code a Small Model transient assembly language program.

```

cseg
.
.   (code)
dseg
org  100h
.
.   (data)
end

```

4.1.3 The Compact Memory Model

The Compact Model is assumed when code and data groups are present, along with one or more of the remaining stack, extra, or auxiliary groups. In this case, the P_CLI system call sets the CS, DS, and ES registers to the base addresses of their respective areas, with the IP set to 0000H, and the SS and SP registers set to a 96-byte stack area allocated by this system call.

Figure 4-4 shows the initial configuration of the segments in the Compact Model. The values of the various segment registers can be changed during execution by loading from the initial values placed in Base Page. This allows access to the entire memory space.




Figure 4-4. Concurrent CP/M Compact Memory Model

If the assembly language transient program intends to use the stack group as a stack area, the SS and SP registers must be set upon entry. The SS and SP registers remain in the initial stack area, even if a stack group is defined.

Although it appears that the SS and SP registers should be set to address the stack group, there are two contradictions. First, the assembly language transient program might be using the stack group as a data area. In that case, the stack values set by the P_CLI system call to allow a far return to terminate a transient program could overwrite data in the stack area. Second, the SS register would logically be set to the base of the group, while the SP would be set to the offset of the end of the group. However, if the stack group exceeds 64K, the address range from the base to the end of the group exceeds a 16-bit offset value.

The following ASM-86 example shows how to code a Compact Model assembly language transient program.

```

cseg
.
.   (code)
dseg
org 100h
.
.   (data)
eseg
.
.   (more data)
sseg
.
.   (stack area)
end

```

4.2 GENCMD

The GENCMD utility creates a CMD file from an input H86 file. GENCMD does not alter the original H86 file. The GENCMD invocation has the following form:

```
GENCMD filename {parameter-list}
```

where the filename corresponds to the H86 input file with an assumed and unspecified filetype of H86. GENCMD accepts optional parameters to specifically identify the 8080 Model and to describe memory requirements of each segment group. The GENCMD parameters are listed following the filename, as shown in the command line above where the parameter list consists of a sequence of keywords (shown below) and values separated by commas or blanks.

```
8080  CODE  DATA  EXTRA  STACK  X1  X2  X3  X4
```

The 8080 keyword forces a single code group so that the P_LOAD system call sets up the 8080 Model for execution, allowing intermixed code and data in a single segment. The form of this command is

```
GENCMD filename 8080
```

The remaining keywords follow the filename or the 8080 option and define specific memory requirements for each segment group, corresponding one-to-one with the segment groups defined in the previous section. In each case, the values corresponding to each group are enclosed in square brackets and separated by commas. Each value is a hexadecimal number representing a paragraph address or segment length in paragraph units denoted by hhhh, prefixed by a single letter that defines each value:

Ahhhh	Load the group at absolute location hhhh
Bhhhh	The group starts at hhhh in the hex file
Mhhhh	The group requires a minimum of hhhh * 16 bytes
Xhhhh	The group can address a maximum of hhhh * 16 bytes

Generally, the CMD file header record values are derived directly from the H86 file and the parameters shown above need not be included. The following situations, however, require the use of GENCMD parameters.

- The 8080 keyword is included whenever ASM-86 is used in the conversion of 8080 programs to the 8086/8088 environment when code and data are intermixed within a single 64K segment, regardless of the use of CSEG and DSEG directives in the source program.
- An absolute address (a hexadecimal value) must be given for any group that must be located at an absolute location. This value is not usually specified, as Concurrent CP/M cannot ensure that the required memory region is available. In that case the CMD file cannot be loaded.
- The B value is used when GENCMD processes a HEX file produced by Intel's OH86 or a similar utility program that contains more than one group. The output from OH86 consists of a sequence of data records with no information to identify code, data, extra, stack, or auxiliary groups. In this case, the B value marks the beginning address of the group named by the keyword, causing GENCMD to load data following this address to the named group (refer to the examples below). Thus, the B value is usually used to mark the boundary between Code and Data Segments when no segment information is included in the HEX file. Files produced by ASM-86 do not require the use of the B value because segment information is included in the H86 file.

- The minimum memory value (M value) is included only when the HEX records do not define the minimum memory requirements for the named group. Generally, the code group size is determined precisely by the data records loaded into the area. The total space required for the group is defined by the range between the lowest and highest data byte addresses. The data group, however, might contain uninitialized storage at the end of the group. Thus no data records are present in the HEX file that define the highest referenced data item. The highest address in the data group can be defined within the source program by including the ASM86 directive DB 0 as the last data item in the assembly language source file. Alternatively, the M value can be included to allocate the additional space at the end of the group. Similarly, the stack, extra, and auxiliary group sizes must be defined using the M value unless the highest addresses within the groups are implicitly defined by data records in the HEX file.
- The maximum memory size, given by the X value, is generally used when additional free memory might be needed for such purposes as I/O buffers or symbol tables. If the data area size is fixed, then the X parameter need not be included. In this case, the X value is assumed to be the same as the M value. The value XFFFF allocates the largest memory region available but, if used, the assembly language transient program must be aware that a three-byte length field is produced in the Base Page for this group where the high-order byte might be nonzero. Programs converted directly from an 8080 environment or programs that use a 2-byte pointer to address buffers should restrict this value to XFFF or less, producing a maximum allocation length of OFFF0H bytes.

The following GENCMD command line transforms the file X.H86 into the file X.COM with the proper header record:

```
A>GENCMD x code[a40] data[m30,xfff]
```

In this case, the code group is forced to paragraph address 40H or its equivalent, byte address 400H. The data group requires a minimum of 300H bytes, but can use up to OFFF0H bytes, if available.

Assuming a file Y.H86 exists on drive B containing Intel HEX records with no interspersed segment information, the command

```
A>GENCMD b:y data[b30,m20] extra[b50] stack[m40] x1[m40]
```

produces the file Y.COM on drive B by selecting records beginning at address 0000H and less than 0300H for the Code Segment, with records starting at 0300H and less than 0500H allocated to the Data Segment. The Extra Segment is filled from records beginning at 0500H and higher, while the Stack and Auxiliary Segment #1 are uninitialized areas requiring a minimum of 0400H bytes each. In this example, the data area requires a minimum of 0200H bytes. Note again that the B value need not be included if the Digital Research ASM-86 assembler is used.


4.3 Intel Hexadecimal File Format

GENCMD input must be in Intel hexadecimal file format, produced by both the Digital Research ASM-86 assembler and the standard Intel OH86 utility program. (Refer to *Intel MCS-86 Software Development Utilities Operating Instructions for ISIS-IISM Users*, published by Intel.) The CMD file produced by GENCMD contains a header record defining the memory model and memory size requirements for loading and executing the CMD file.


An Intel hexadecimal file consists of the traditional sequence of ASCII records where the beginning of the record is marked by an ASCII colon, and each subsequent digit position contains an ASCII hexadecimal digit in the range 0-9 or A-F.

There are four kinds of hexadecimal record formats. The Start Address Record specifies the starting address of the execution file. The Extended Address Record specifies the bits 4-19 of the Segment Base Address, where bits 0-3 of the SBA are zero. The Data Record contains a string of hexadecimal ASCII code that represents a portion of the 8086 memory image. The End-of-File record specifies the end of the object file.


Figure 4-5 shows the four record formats, their fields, and the contents of these fields. The fields are defined in Table 4-2.




STARTING ADDRESS RECORD



EXTENDED ADDRESS RECORD



DATA RECORD



END OF FILE RECORD

Figure 4-5. Intel Hexadecimal File Formats

Table 4-2. Intel Hex Field Definitions

<i>Field</i>	<i>Contents</i>
Rec Mark	Specifies start of record
Rec Len	Record Length 00-FF (0-255 in decimal)
zeros	Extended Address Record: 0000H Starting Address Record: 0000H End-of-File Record: 0000H
Ld Addr	Data Record: SBA offset defining address of byte 0 of data
Rec Type	00 = Data Record 01 = End-of-File Record 02 = Extended Address Record 03 = Starting Address Record The following are output from ASM-86 only: 81 same as 00, data belongs to Code Segment 82 same as 00, data belongs to Data Segment 83 same as 00, data belongs to Stack Segment 84 same as 00, data belongs to Extra Segment * 85 paragraph address for absolute Code Segment * 86 paragraph address for absolute Data Segment * 87 paragraph address for absolute Stack Segment * 88 paragraph address for absolute Extra Segment
C-Seg	Four hexadecimal digits specifying the Code Segment address. The high-order and low-order digits are the 10th and 13th characters of the record, respectively.
USBA	Four hexadecimal digits specifying the Upper Segment Base Address. The high-order and low-order digits are the 10th and 13th characters of the record, respectively.
data	Pairs of hexadecimal digits representing the ASCII code for each data byte. The high-order digit is the first digit of each pair.

Table 4-2. (continued)

<i>Field</i>	<i>Contents</i>
Checksum	<p>Extended Address Record: Checksum of Rec Len, zeros, Rec Type, and USBA fields.</p> <p>Starting Address Record: Checksum of Rec Len, zeros, Rec Type, C-Seg, and IP fields.</p> <p>Data Record: Checksum of Rec Len, Ld Addr, Rec Type, and data fields.</p> <p>End-of-File Record: Contains ASCII code 4646H, checksum of Rec Len, zeros, and Rec Type fields.</p>
* 85, 86, 87, and 88 are Digital Research Extensions.	

All characters preceding the colon for each record are ignored. See *MCS[®]-86 Absolute Object File Formats*, published by Intel, for additional information on hexadecimal file record format.

End of Section 4

Section 5

Resident System Process Generation

5.1 Introduction to RSPs

Resident System Processes are programs that become part of the Concurrent CP/M operating system. They can be useful in several ways: to create a turnkey system, autoloading programs when Concurrent CP/M is booted; to build customized user interfaces or shells at the consoles, for monitoring hardware not supported in the XIOS; and to avoid disk loading time for frequently-used commands.

The source code for the ECHO RSP is included in Appendix D. Study this listing carefully while reading this section. The discussion of the P_CREATE system call in Section 6 is also helpful in understanding RSPs.

Resident System Processes are included in Concurrent CP/M during system generation. GENCCPM searches the directory for all files with the filetype RSP and prompts the user to choose whether it is to be included in the generated system file, CCPM.SYS. An RSP file is created by generating a CMD file and renaming it with an RSP filetype. The GENCCPM program is documented in the *Concurrent CP/M System Guide*.

5.2 RSP Memory Models

Under Concurrent CP/M, there are two basic memory models for RSPs. They are similar to the 8080 Model and the Small Model of transient programs. However, several important distinctions exist between the transient program and RSP memory models. The RSP has no equivalent to the Base Page of the transient program's Data Segment. The RSP is responsible for its own Process Descriptor (PD) and User Data Area (UDA). The RSP must also allocate an additional 96 bytes at the end of the User Data Area if 8087 processing is required. The system creates and initializes these data structures for the transient programs automatically at load time. RSPs, on the other hand, must initialize these structures within their own Data Segments (See P_CLI and P_CREATE system calls for PD and UDA descriptions).

Note that Concurrent CP/M does not support compact model RSPs. Extra and Stack Segments must be part of the Data Segment.

Although there is no Base Page in an RSP, there is an RSP header that must exist at offset 00H of the Data Segment. In the 8080 Model, this implies that the RSP header is in the Code Segment. The RSP header and the associated data structures are discussed in Section 5.4.

5.2.1 8080 Model RSP

The 8080 Model consists of mixed code and data. When the system gives control of the CPU to an 8080 Model RSP, it initializes the Code, Data, Extra and Stack Segment registers to the same value. Use GENCMD with the 8080 option to generate an 8080 Model RSP. GENCCPM assumes the 8080 Model if the CMD File Header Record of the RSP has a single Code Group Descriptor and no other Group Descriptors (refer to Section 3.2). When discussing an 8080 Model RSP, any reference to the Data Segment also refers to the Code Segment.

5.2.2 Small Model RSP

The Small Model RSP implies separate Code and Data Segments. Before the system gives control of the CPU to a Small Model RSP, it initializes the Data, Extra and Stack Segment Registers to the Data Segment address, while the Code Segment register is initialized to the Code Segment address. There is no guarantee where GENCCPM will place the Code Segment in memory relative to the Data Segment. The CMD Header Record for this kind of RSP must have both Data and Code Group Descriptors.




Figure 5-1. 8080 and Small RSP Models

5.3 Multiple Copies of RSPs

At system generation, GENCCPM can make up to 255 extra copies of an RSP, such that each copy generates a separate process running under Concurrent CP/M. GENCCPM accomplishes this by making multiple copies of the RSP, and initializing each to be a separate RSP. The number of copies made by GENCCPM can be fixed, or dependent on a byte value in the System Data Area. To determine the number of copies to make, GENCCPM looks at two fields in the RSP Header. The format of the RSP Header is shown in Figure 5-2.




Figure 5-2. RSP Header Format

If the SDATVAR field is nonzero, it is used as an offset of a byte value in the System Data Area, which contains the number of copies to be generated. The offset should indicate a value that is set by the user during GENCCPM. The TMP RSP uses this feature by placing the offset of the NVCNS (Number of Virtual Consoles) field into the SDATVAR field. This way, a TMP is generated for each System Console specified by the user. If SDATVAR is 0 then the NCP byte in the RSP header is used as the number of extra copies to make. If both of these fields in the RSP Header are 0 then no extra copies are made, and only a single RSP is created. The ECHO RSP is an example of the latter.

If the number of extra copies is determined by GENCCPM to be greater than 0, each copy of the RSP is given a unique copy number. The copy number is placed in the NCP field and the ASCII equivalent is appended to the end of the Process Descriptor NAME field of each copy. If there is not enough space for the number in the PD NAME, part of the PD NAME is over written. For the example TMP RSP, GENCCPM makes the specified number of copies and changes the NAME field in each copy to be TMP0, TMP1, TMP2,..., and sets the NCP field to 0, 1, 2,..., respectively.

5.3.1 8080 Model

When GENCCPM makes copies of an 8080 Model RSP, the CS, DS, ES, and SS fields in each copy's User Data Area are set to the paragraph address where the RSP is in memory after loading.

5.3.2 Small Model

If multiple copies of a Small Model RSP are to be generated, GENCCPM copies both the Code and Data Groups of the RSP, if the MEM field of the Process Descriptor is 0. See the P_CREATE system call for a description of the Process Descriptor format. GENCCPM sets the UDA fields CS to the Code Segment of the RSP and DS, ES and SS to the Data Segment of the RSP.

5.3.3 Small Model with Shared Code

If a Small Model RSP has a nonzero MEM field in its Process Descriptor, the Code Segment is assumed to be reentrant. When copies are made of this type of RSP only the Data Group is copied. GENCCPM sets the UDA CS field for each copy to the paragraph address of the one Code Segment for the RSP's. The DS, ES, and SS, in each copied Data Segment, are set by GENCCPM to the paragraph address of the Data Segment for that particular copy.

5.4 Creating and Initializing an RSP

An RSP that is to be invoked from a console, or through the P_CLI system call, must create a special queue called an RSP Command Queue. Such an RSP is called a Command RSP. This type of RSP usually performs some initialization routine and then goes into a loop. The initialization routine consists of creating and opening an RSP Command Queue as well as changing the priority to the default transient process priority. (Priority values with regard to RSPs are discussed below.)

The first step of the loop reads a message from the RSP Command Queue. The process that writes the message to the RSP Command Queue activates the associated RSP. After the RSP returns from the Q_READ system call, it obtains the system resources it needs, such as the calling process' console. Typically, the RSP process is assigned the console process by the CLI after the CLI has succeeded in writing the command tail to the RSP Queue. This is only true if the RSP Process Descriptor name matches the RSP Command Queue name. Refer to the P_CLI (Call Command Line Interpreter) system call description for information about how the CLI handles a command.

When the RSP completes its activities for the given command, it releases any system resources it has acquired, including the console, and restarts the loop by reading from its RSP Command Queue. A Command RSP is a single process and is a serially reusable resource; in other words, the RSP acts on one message at a time. When several processes attempt to invoke a single Command RSP, they wait as described in the Q_READ and Q_CREAD system call in Section 6. Refer to these and to the Q_WRITE and Q_CWRITE system calls for further details.

Note: It is certainly possible to create RSPs that are invoked differently.

The format of the RSP Command Queue Message is shown in Figure 5-3.




Figure 5-3. RSP Command Queue Message

The PDADDRESS is the offset relative to the System Data Area segment of the Process Descriptor of the process calling the RSP. A program that wants to invoke an RSP and is forming an RSP Command Queue Message, can find its Process Descriptor address by calling the P_PDADR system call. The COMMAND TAIL usually contains what the TMP sends to the CLI minus the command name, and is terminated with a zero byte.

When a command is entered at a console, the TMP performs a P_CLI system call. The P_CLI system call attempts to open a queue that has the RSP Flag on and has the same name as the command sent to the CLI. If the Q_OPEN is successful, the P_CLI system call attempts to assign the calling process's console to a process with the same name as the command. The P_CLI system call then creates an RSP Command Queue Message with the command tail sent to the CLI from the TMP, and writes it to the RSP Command Queue (refer to the discussion of the P_CLI and Q_WRITE system calls in Section 6). A transient program can use a Command RSP in the same manner by writing directly to the appropriate RSP Command Queue. An advantage of using the P_CLI system call is that it looks for an RSP first and only searches on disk for a CMD file if the the RSP is not found.

When an RSP reads an RSP Command Queue Message, it often needs information about the calling process, such as which console, list device, drive, or user number to use. If an RSP is invoked through the P_CLI system call, the RSP is assigned the calling process's console, but if the RSP Command Queue is written to directly, the calling process might or might not assign its console to the RSP. A Command RSP can use the PD address in the Command RSP Message to find out what the default devices of the calling process are. The RSP should release any resources it assigns to itself when it is finished.

The beginning of the RSP Data Segment has a fixed format starting at offset 0. This data structure is the RSP Header. Note that in the 8080 Model, the RSP Header is also in the Code Segment. After the RSP Header is a Process Descriptor starting at offset 010H. A User Data Area and a stack must also be within the Data Segment, with the UDA placed at a paragraph boundary relative to the beginning of the Data Segment. If system calls assuming a default DMA buffer are used, a 128-byte DMA Buffer must also exist. The DMA OFFSET field in the User Data Area should be set to the address of the DMA buffer. When the process is created by Concurrent CP/M, the DMA SEGMENT field is initialized to the same value as the DS register. The DMA SEGMENT and OFFSET can also be set by calling F_DMASEG and F_DMAOFF once the RSP is running. The beginning of the RSP Data Segment is shown in Figure 5-4.




Figure 5-4. RSP Data Segment

The RSP Header must be located at offset zero in the RSP Data Segment, the RSP Process Descriptor must be at offset 010H, and the RSP User Data Area must begin on an even paragraph boundary.

5.4.1 The RSP Header

As discussed in Section 5.2, the number of copies made of an RSP is dependent on the values of the SDATVAR and NCP fields in the RSP Header. If no copies are desired, these fields must be zero. As a convenience, when Concurrent CP/M creates the RSP process, the LINK field in the RSP Header is set to the paragraph address of the System Data Area. The System Data Area can always be obtained by an RSP or transient program with the S_SYSDAT system call.

5.4.2 The RSP Process Descriptor

The RSP Process Descriptor should be initialed to zeros, except for the PRIORITY, FLAGS, NAME, and UDA SEGMENT fields. The PRIORITY field is usually initialized to 190. This is higher than transient programs and TMPs (200 and 198 respectively), but lower than the INIT process, which has a priority of 1. The description of the P_PRIORITY system call in Section 6 contains more information about system priority assignments.

Starting an RSP at a priority of 190 ensures that the RSP is able to create and open an RSP Command Queue before it can be invoked through a TMP. RSPs such as ECHO usually set their priority to 200 after creating and opening their RSP Command Queue and before attempting to read from the queue.

Note: There are no guarantees about the order in which the RSP processes are created by the Concurrent CP/M operating system. If one RSP must run before another, it must have a higher priority. Such is the case when one RSP uses a resource created by a second RSP; the second must run (at least during initialization) with a priority higher than the first.

The Process Descriptor SYS and KEEP Flags can be initialized in the RSP Data Segment (refer to P_CREATE in Section 6 for further flag details). The SYS Flag allows a process to read and write to and from restricted system queues. This is discussed below with regard to RSP Command Queues. The KEEP flag signals to the operating system that this process cannot be terminated. This flag is necessary if an RSP is not to be terminated when a CTRL-C is typed on a console being used by the RSP. The 8087 flag tells the system that a process is actively using the 8087 processor.

The NAME field of the RSP's Process Descriptor is 8 bytes long. It is assumed to be left-justified and padded with blanks on the right. If an RSP Command Queue is going to be used to invoke the RSP through the CLI, the PD must have the same uppercase name as the Command Queue. The UDA field in the Process Descriptor must be the offset in paragraphs of the UDA relative to the RSP data segment. GENCCPM restores the UDA field in the Process Descriptor to the actual UDA paragraph address when the system is generated.

If the PD field name is not the same as the Command Queue, the console is not assigned to the RSP by the CLI.

5.4.3 The RSP User Data Area

The User Data Area must have its SP field set to the offset of a three-word IRET structure, in the RSP's Data Segment. The offset is relative to the beginning of the Data Segment. The first of the three words is the offset of the code entry point for the RSP, relative to the beginning of the RSP Code Segment. Concurrent CP/M executes an IRET instruction to start the RSP using these three words for the IP, CS and Flag registers respectively. The CS value on the stack is initialized to be the CS field of the UDA, while the Flag value is set to 0200H (interrupts on). The RSP stack must come immediately before these three words.

The initial values of the AX, BX, CX, DX, DI, SI, and BP registers are taken from the appropriate fields in the UDA.

The DMA OFFSET field should be set to the offset of the DMA buffer in the RSP's Data Segment. Except for the SP and DMA OFFSET fields, and possibly the AX, BX, CX, DX, DI, SI, and BP fields, the remainder of the UDA fields should be initialized to 0. The CS, DS, ES, and SS fields are set by GENCCPM as discussed in Section 5.3.

If you include the 8087 extension in the UDA, you must initialize the CW field (Control Word) to 03FFH and the SW (Status Word) field to 0 before system generation.

5.4.4 The RSP Stack

The RSP must reserve space for its stack, which is assumed to lie within the RSP's Data Segment. This stack must be large enough to accommodate what the RSP code needs, plus four levels (eight bytes) to handle possible hardware interrupts. We highly recommend that you reserve more than four extra levels of stack.

The SP field in the RSP's UDA points to the top of this stack; the top contains the three-word IRET instruction discussed above.

5.4.5 The RSP Command Queue

The RSP's Command Queue contains information that determines when it begins execution, and to which console it is attached. If an RSP is to be accessible from a console via the TMP, the Command Queue name must be in uppercase. The FLAGS field in the RSP Command Queue Descriptor must have the RSP bit on. If this flag is not on, the CLI will not write a message to the RSP Command Queue, and instead attempts to load a transient program. The KEEP flag should be set on to protect the RSP QUEUE from inadvertent use of the Q_DELETE system call.

The RESTRICTED flag (refer to the Q_MAKE system call in Section 6) makes a queue accessible only by privileged processes. Privileged processes have the SYS Flag on in their Process Descriptor. If the RESTRICTED Flag is on in an RSP Command Queue, then only privileged processes can invoke the related RSP. A lowercase letter in the RSP Command Queue name and the RESTRICTED Flag provide two methods of filtering access to an RSP QUEUE.

The Queue Descriptor of the RSP Command Queue must have a message length of 131 bytes. The format of this message is shown above. The number of messages is usually 1. If the Queue Descriptor is within 64K bytes of the beginning of the System Data Area, buffer space for the Queue Descriptor must be allocated in the RSP. The BUFFER field in the Queue Descriptor must be the offset of this buffer, relative to the beginning of the RSP's Data Segment. The buffer size is the message length times the number of messages, usually 131 bytes.

Note: The queue buffer should be before the Queue Descriptor within the RSP Data Segment.

An RSP can certainly create other queues besides the RSP Command Queue used with Command RSPs. However, any queue an RSP creates that lies within 64K of the System Data Area must have a buffer area pointed to by the BUFFER field in its Queue Descriptor. To be safe, the buffer should come before the Queue Descriptor in the RSP's Data Segment. It is assumed the BUFFER field points to a buffer that is also within 64K of the System Data Area. If the Queue Descriptor is farther than 64K from the System Data Area, Concurrent CP/M uses buffer space in the System Data Area. Refer to the Q_MAKE system call in Section 6 for further details.

In order to open the RSP Command Queue and subsequently read from it, a Queue Parameter Block and its associated buffer must be allocated in the RSP's Data Segment. These structures are treated just as in a transient process. For any queues created by an RSP, it is stressed that the queue buffer areas associated with the Queue Descriptor and the Queue Parameter Block are separate, distinct areas of storage.

5.4.6 Multiple Processes within an RSP

An RSP can create child processes by calling the P_CREATE system call. Note that if the Process Descriptor of the process being created is within 64K bytes of the beginning of the System Data Area, the PD structure is used directly by Concurrent CP/M. Otherwise the PD structure is copied into the PD table in the System Data Area.

5.5 Developing and Debugging an RSP

The first RSP you attempt should be very simple, on the order of complexity of the ECHO RSP listed in Appendix D. New RSPs should be developed and debugged as if they were transient processes, such as Concurrent CP/M CMD utilities, then converted into RSPs.

An RSP debugging session should proceed like an XIOS debugging session: first load CP/M-86, then invoke DDT-86[™], and then bring up Concurrent CP/M. The *Concurrent CP/M System Guide* provides more information about running Concurrent CP/M under CP/M-86.

After reading in the CCPM.SYS file under DDT-86, find the RSPSEG field of the System Data Segment (SYSDAT). The paragraph address of the SYSDAT is found in the A_BASE field of the Data Group Descriptor in the CCPM.SYS command file header. The CMD header is described in Section 3.2 and the SYSDAT area is described in the S_SYSDAT system call in Section 6. The RSPSEG field contains the paragraph address of the Data Segment of the first RSP in a linked list of the RSPs included by GENCCPM.

By using the Display Memory (D) command of DDT-86 to show memory at the segment RSPSEG, the name of the first RSP can be identified in the RSP's Process Descriptor. The LINK field in the RSP Header, which will be the first word in the RSPSEG segment, is the paragraph value of the next RSP's Data Segment. A zero in the LINK field means the end of the list of RSPs. Note that linkage information is lost once Concurrent CP/M is initialized. The LINK field of the RSP Header contains the System Data Segment once an RSP begins execution.

Once the RSP to be debugged is located, the initial code entry point can also be found. As discussed previously, the SP field in the RSP's UDA is the offset from the beginning of the RSP's Data Segment of the three-word IRET structure. The first word of the IRET structure contains the initial value of the IP register when Concurrent CP/M creates the RSP process. The initial value of the CS register is in the CS field also in the RSP's UDA. Once this is done, you can set break points in the RSP, similar to setting break points in XIOS system calls.

End of Section 5

Section 6

System Calls

This section describes the Concurrent CP/M system calls in tabular form. It is intended both as an introduction to the calls and as a reference for use during programming. You should be familiar with the material in Sections 1 through 5 before proceeding.

The first table, Table 6-1, describes the categories of Concurrent CP/M system calls and their general uses. Table 6-2 summarizes the Concurrent CP/M system calls. Use it as a quick reference to find the system call you need while programming. The system calls are broken down into functional groups. Immediately following is Table 6-3, a cross-reference showing the system calls in numerical order. Table 6-4 is an index providing the page numbers and figure titles of commonly used data structures. Table 6-5 lists the error codes returned in register CX.

Table 6-1. System Call Categories

<i>Category</i>	<i>Use</i>
C_ Console System Calls	The Console System Calls handle I/O operations for virtual consoles on a character, string, and line basis, attach and detach consoles from processes, and return or change the number corresponding to the default virtual console.
DEV_ Device System Calls	The Device System Calls deal with flags and polling in managing system resources.
DRV_ Disk Drive System Calls	The Disk Drive System Calls manage Concurrent CP/M logical drives.
F_ File-Access System Calls	The File-Access System Calls include calls that operate on files within a directory, calls that operate on records within files, and miscellaneous system calls related to file I/O.
L_ List Device System Calls	The List Device System Calls write characters or strings to the default list device, attach and detach the default list device from calling processes, and return or change the number corresponding to the default list device.
M_ MP/M-86 [™] Memory Management System Calls	The M_ Memory Management System Calls are included for compatibility with MP/M-86. These calls allocate and free memory segments according to the MP/M-86 segmentation algorithm.

Table 6-1. (continued)

<i>Category</i>	<i>Use</i>
MC_ CP/M-86 Memory Management System Calls	The MC_ Memory Management System Calls allocate and free memory segments according to the CP/M-86 segmentation algorithm.
P_ Process/Program System Calls	The Process/Program System Calls create and terminate processes, call other processes, and perform other operations on processes.
Q_ Queue Management System Calls	The Queue Management System Calls create, delete, open, read from, and write to queues.
S_ System Calls	The System Calls return various types of systems data, such as version numbers and addresses.
T_ Time System Calls	The Time System Calls set the system calendar and clock and return the time from them in hours and minutes or in hours, minutes, and seconds.

Table 6-2. Concurrent CP/M System Calls

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
Console I/O System Calls			
149	95	C_ASSIGN	Assign default virtual console to another process.
146	92	C_ATTACH	Establish ownership of the default virtual console to the calling process; suspend process until console becomes available.
162	A2	C_CATTACH	Conditionally establish ownership of the default virtual console by the calling process; return an error message if the device is unavailable.
110	6E	C_DELIMIT	Set or return current String Output Delimiter. Used with C_WRITESTR.
147	93	C_DETACH	Detach default virtual console from the calling process.
153	99	C_GET	Return the virtual console number of the calling process.
109	6D	C_MODE	Set or return Console mode.
6	06	C_RAWIO	Perform Raw mode I/O with the default virtual console.
1	01	C_READ	Read a character from the default virtual console.
10	0A	C_READSTR	Read an edited line from the default virtual console.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
148	94	C_SET	Set or change the default virtual console for the calling process.
11	0B	C_STAT	Obtain the input status of the default virtual console.
2	02	C_WRITE	Write a character to the default virtual console.
111	6F	C_WRITEBLK	Write a specified number (block) of characters to the default virtual console.
9	09	C_WRITESTR	Write a string to the default virtual console until delimiter.
Device System Calls			
133	85	DEV_SETFLAG	Set a system flag.
132	84	DEV_WAITFLAG	Wait for a system flag to be set before restoring the current process.
131	83	DEV_POLL	Poll a noninterrupt-driven device.
Disk Drive System Calls			
38	26	DRV_ACCESS	Indicate access to specified drives.
27	1B	DRV_ALLOCVEC	Get the address of the disk Allocation Vector.
13	0D	DRV_ALLRESET	Reset all disk drives.
31	1F	DRV_DPB	Return the segment and offset address of the Disk Parameter Block for the default disk of the calling process.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
48	30	DRV_FLUSH	Write internal pending blocking/deblocking data buffers to disk.
39	27	DRV_FREE	Relinquish access to specified drives.
25	19	DRV_GET	Return the default drive of the calling process.
101	65	DRV_GETLABEL	Return the directory label data byte for the specified drive.
24	18	DRV_LOGINVEC	Return bit map of logged-in disk drives.
37	25	DRV_RESET	Reset the specified drives.
29	1D	DRV_ROVEC	Return bit map vector of drives set to Read-Only.
14	0E	DRV_SET	Set default drive of calling process.
100	64	DRV_SETLABEL	Create or update a directory label.
28	1C	DRV_SETRO	Set the default drive to Read-Only.
46	2E	DRV_SPACE	Return unallocated space on the specified drive.
Disk File System Calls			
30	1E	F_ATTRIB	Set file attributes.
16	10	F_CLOSE	Close file.
19	13	F_DELETE	Delete file.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
52	34	F_DMAGET	Return segment and offset address of Direct Memory Address buffer.
26	1A	F_DMAOFF	Set the Direct Memory Address offset address.
51	33	F_DMASEG	Set Direct Memory Address buffer segment address.
45	2D	F_ERRMODE	Set the BDOS Error mode.
42	2A	F_LOCK	Lock record within file opened in Unlocked mode.
22	16	F_MAKE	Create file.
44	2C	F_MULTISEC	Set the BDOS Multisector Count.
15	0F	F_OPEN	Open file for record access.
152	98	F_PARSE	Parse an ASCII string and initialize an FCB.
106	6A	F_PASSWD	Set the default password.
36	24	F_RANDREC	Set the Random Record field in the FCB from the sequential record position.
20	14	F_READ	Read record sequentially.
33	21	F_READRAND	Read random record.
23	17	F_RENAME	Rename file.
17	11	F_SFIRST	Search for first matching directory FCB that matches the specified FCB.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
35	23	F_SIZE	Return the size of a file.
18	12	F_SNEXT	Search for next matching directory FCB that matches the FCB specified in the F_SFIRST system call.
102	66	F_TIMEDATE	Return file's date and time stamps and password mode.
99	63	F_TRUNCATE	Truncate file to the specified Random Record Number.
43	2B	F_UNLOCK	Remove record locks.
32	20	F_USERNUM	Set or return the default user number of the calling process.
21	15	F_WRITE	Write records sequentially.
34	22	F_WRITERAND	Write random records.
103	67	F_WRITEXFCB	Create or update file's XFCB.
40	28	F_WRITEZF	Write random records and zero-fill any previously unallocated data blocks.
List Device System Calls			
158	9E	L_ATTACH	Establish ownership of the default list device by the calling process; suspend the process until the device is available.
161	A1	L_CATTACH	Conditionally establish ownership of the default list device by the calling process; return error code if the device is unavailable.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
159	9F	L_DETACH	Relinquish ownership of the default list device.
164	A4	L_GET	Return the default list device number of the calling process.
160	A0	L_SET	Change the default list device for the calling process.
5	05	L_WRITE	Write a character to the default list device.
112	70	L_WRITEBLK	Write the specified number of characters (block) to the default list device.
MP/M Compatible Memory Allocation System Calls			
128	80	M_ALLOC	Allocate the memory segment between the sizes specified in the Memory Parameter Block to the calling process.
129	81	same as 128	
130	82	M_FREE	Free the specified memory segment.
CP/M Compatible Memory Allocation System Calls			
54	36	MC_ABS	Allocate the maximum amount of RAM available at a specified address.
58	3A	MC_ALLFREE	Free all memory owned by the calling process.
55	37	MC_ALLOC	Allocate a segment of RAM, as specified in the Memory Control Block, to the calling process.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
56	38	MC_ALLOCABS	Allocate a specified amount of RAM, as above, but beginning at a specific address.
57	39	MC_FREE	Free an area of RAM beginning at a specified address, and extending to the end of the previously-allocated memory area.
53	35	MC_MAX	Allocate the maximum amount of RAM available in the system.
Process/Program System Calls			
157	9D	P_ABORT	Terminate a process specified by name or Process Descriptor address.
47	2F	P_CHAIN	Load, initialize, and jump to the program specified in the DMA buffer.
150	96	P_CLI	Interpret and execute the specified command line by calling Command Line Interpreter (CLI).
144	90	P_CREATE	Create a subprocess.
141	8D	P_DELAY	Suspend the calling process for a specified number of system clock ticks.
142	8E	P_DISPATCH	Force a dispatch operation; give up the CPU resource to the highest priority process ready to run.
59	3B	P_LOAD	Load the specified CMD file in memory; return its base page segment address.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
156	9C	P_PDADR	Return the address of the Process Descriptor of the calling process.
145	91	P_PRIORITY	Set the priority of the calling process.
151	97	P_RPL	Invoke a system call from a Resident Procedure Library.
143	8F	P_TERM	Terminate the calling process.
0	00	P_TERMCPM	Terminate calling process unconditionally, release all owned resources.
Queue System Calls			
138	8A	Q_CREAD	Conditionally read a message from a system queue; return error code if a message is not available.
140	8C	Q_CWRITE	Conditionally write a message to a system queue; return an error code if space is not available.
136	88	Q_DELETE	Delete a system queue.
134	86	Q_MAKE	Create a system queue.
135	87	Q_OPEN	Open a system queue for subsequent queue operations.
137	89	Q_READ	Read a message from a system queue; suspend calling process until message is available.

Table 6-2. (continued)

<i>Number</i>		<i>Mnemonic</i>	<i>Definition</i>
<i>Dec</i>	<i>Hex</i>		
139	8B	Q_WRITE	Write a message to a system queue; suspend calling process until space becomes available.
System System Calls			
12	0C	S_BDOSVER	Return BDOS version number, CPU and operating system type.
50	32	S_BIOS	Call specified CP/M-86 BIOS character I/O routine.
163	A3	S_OSVER	Return type and version number of Concurrent CP/M.
107	6B	S_SERIAL	Return the Concurrent CP/M system serial number.
154	9A	S_SYSDAT	Return address of the System Data Segment (Sysdat)
Time System Calls			
105	69	T_GET	Obtain the system calendar and clock, hours and minutes only.
155	9B	T_SECONDS	Return current system date and time; hours, minutes, seconds.
104	68	T_SET	Set internal system calendar and clock to specified value.

6.1 System Call Summary

Table 6-3 lists the Concurrent CP/M system calls in summary form, including the parameters a process must pass when calling the system call, and the values the system returns to the process.

Appendix A lists the Concurrent CP/M system calls by function number, and includes all the information in Table 6-3.

Table 6-3. System Call Summary

<i>Mnemonic</i>	<i>Dec</i>	<i>Hex</i>	<i>Input Parameters</i>	<i>Returned Values</i>
C_ASSIGN	149	95	DX = .ACB	AX = Rtn Code
C_ATTACH	146	92	none	none
C_CATTACH	162	A2	none	AX = Rtn Code
C_DELIMIT	110	6E	DX = Out Delim	AL = Out Delim
C_DETACH	147	93	none	none
C_GET	153	99	none	AL = con #
C_MODE	109	6D	DX = Con Mode = 0FFFFH	none AX = Con Mode
C_SET	148	94	DL = Console	none
C_RAWIO	6	6	see def	see def
C_READ	1	1	none	AL = char
C_READSTR	10	A	DX = .Buffer	see def
C_STAT	11	B	none	AL = 00/01
C_WRITE	2	2	DL = char	none
C_WRITEBLK	111	6F	DX = .CHCB	none
C_WRITESTR	9	9	DX = .Buffer	none
DEV_POLL	131	83	DL = Device	none
DEV_SETFLAG	133	85	DL = Flag	AX = Rtn Code
DEV_WAITFLAG	132	84	DL = Flag	AX = Rtn Code
DRV_ACCESS	38	26	DX = drive Vect	none
DRV_ALLOCVEC	27	1B	none	AX = .Alloc
DRV_ALLRESET	13	D	none	see def
DRV_DPB	31	1F	none	AX = .DPB
DRV_FLUSH	48	30	none	see def

Table 6-3. (continued)

<i>Mnemonic</i>	<i>Dec</i>	<i>Hex</i>	<i>Input Parameters</i>	<i>Returned Values</i>
DRV_FREE	39	27	DX = drive Vect	none
DRV_GET	25	19	none	AL = Cur Drive #
DRV_GETLABEL	101	65	DX = Drive #	AL = Label Data Byte
DRV_LOGINVEC	24	18	none	AX = Login Vect.
DRV_RESET	37	25	DX = drive Vect	AL = Err Code
DRV_ROVEC	29	1D	none	AX = R/O Vect.
DRV_SET	14	E	DL = Drive #	see def
DRV_SETLABEL	100	64	DX = .FCB	AL = Dir Code
DRV_SETRO	28	1C	none	see def
F_ATTRIB	30	1E	DX = .FCB	see def
F_CLOSE	16	10	DX = .FCB	AL = Dir Code
F_DELETE	19	13	DX = .FCB	AL = Dir Code
F_DMAGET	52	34	none	AX = DMA Offset
F_DMAOFF	26	1A	DX = .DMA	none
F_DMASEG	51	33	DX = .DMA Seg	none
F_ERRMODE	45	2D	DL = Err Mode	none
F_LOCK	42	2A	DX = .FCB	AL = Err Code
F_MAKE	22	16	DX = .FCB	AL = Dir Code
F_MULTISEC	44	2C	DL = # of Records	AL = Rtn Code
F_OPEN	15	F	DX = .FCB	AL = Dir Code
F_PARSE	152	98	DX = .PFCB	see def
F_PASSWD	106	6A	DX = .Password	none
F_RANDREC	36	24	DX = .FCB	R0, R1, R2
F_READ	20	14	DX = .FCB	AL = Err Code
F_READRAND	33	21	DX = .FCB	AL = Err Code
F_RENAME	23	17	DX = .FCB	AL = Dir Code
F_SFIRST	17	11	DX = .FCB	AL = Dir Code
F_SIZE	35	23	DX = .FCB	R0, R1, R2
F_SNEXT	18	12	none	AL = Dir Code
F_TIMEDATE	102	66	DX = .XFCB	AL = Dir Code
F_TRUNCATE	99	63	DX = .FCB	see def
F_UNLOCK	43	2B	DX = .FCB	AL = Err Code
F_USERNUM	32	20	DL = OFFH (get) = User # (set)	AL = User # none

Table 6-3. (continued)

<i>Mnemonic</i>	<i>Dec</i>	<i>Hex</i>	<i>Input Parameters</i>	<i>Returned Values</i>
F_WRITE	21	15	DX = .FCB	AL = Err Code
F_WRITERAND	34	22	DX = .FCB	AL = Err Code
F_WRITEXFCB	103	67	DX = .XFCB	AL = Dir Code
F_WRITEZF	40	28	DX = .FCB	AL = Err Code
L_ATTACH	158	9E	none	none
L_CATTACH	161	A1	none	AX = Rtn Code
L_DETACH	159	9F	none	none
L_GET	164	A4	none	AL = list #
L_SET	160	A0	DL = List #	none
L_WRITE	5	5	DL = char	none
L_WRITEBLK	112	70	DX = .CHCB	none
M_ALLOC	128	80		
M_ALLOC	129	81	DX = .MPB	AX = Rtn Code
M_FREE	130	82	DX = .MPB	none
MC_ABSALLOC	56	38	DX = .MCB	see def
MC_ABSMAX	54	36	DX = .MCB	see def
MC_ALLFREE	58	3A	none	none
MC_ALLOC	55	37	DX = .MCB	see def
MC_FREE	57	39	DX = .MCB	see def
MC_MAX	53	35	DX = .MCB	see def
P_ABORT	157	9D	DX = .ABP	AX = Rtn Code
P_CHAIN	47	2F	see def	none
P_CLI	150	96	DX = .CLBUF	none
P_CREATE	144	90	DX = .PD	none
P_DELAY	141	8D	DX = #ticks	none
P_DISPATCH	142	8E	none	none
P_LOAD	59	3B	DX = .FCB	AX = BP Addr
P_PDADR	156	9C	none	AX = PD Addr
P_PRIORITY	145	91	DL = Priority	none
P_RPL	151	97	DX = .CPB	AX = result
P_TERM	143	8F	DL = Term.Code	AX = Rtn Code
P_TERMCPM	0	0	none	AX = Rtn Code

Table 6-3. (continued)

<i>Mnemonic</i>	<i>Dec</i>	<i>Hex</i>	<i>Input Parameters</i>	<i>Returned Values</i>
Q_CREAD	138	8A	DX = .QPB	AX = Rtn Code
Q_CWRITE	140	8C	DX = .QPB	AX = Rtn Code
Q_DELETE	136	88	DX = .QPB	AX = Rtn Code
Q_MAKE	134	86	DX = .QD	none
Q_OPEN	135	87	DX = .QPB	AX = Rtn Code
Q_READ	137	89	DX = .QPB	none
Q_WRITE	139	8B	DX = .QPB	none
S_BDOSVER	12	C	none	AX = Version#
S_BIOS	50	32	DX = .BD	AX = BIOS rtn
S_OSVER	163	A3	none	AX = Version #
S_SERIAL	107	6B	DX = .serialnmb	serialnmb set
S_SYSDAT	154	9A	none	AX = Sys Data Addr
T_GET	105	69	DX = .TOD	AL = seconds
T_SECONDS	155	9B	DX = .TOD	TOD filled in
T_SET	104	68	DX = .TOD	none

Note: System calls 3, 4, 7, and 8 are not supported by Concurrent CP/M.

Conventions used in Table 6-3:

.	= Address of	MCB	= Memory Control Block
#	= Number	MPB	= Memory Parameter Block
ACB	= Assign Control Block	Num	= Number
APB	= Abort Parameter Block	Out	= Output
Addr	= Address	PD	= Process Descriptor
BD	= Bios Descriptor	PFCB	= Parse Filename Control Block
BP	= Base Page	QD	= Queue Descriptor
Char	= ASCII Character	QPB	= Queue Parameter Block
CHCB	= Character Control Block	Rec	= Record
CLBUF	= Command Line Buffer	Rtn	= Return
CPB	= Call Parameter Block	Sys	= System
Con	= Console	Term.	= Termination
Cur	= Current	TOD	= Time of Day
Delim	= Delimiter	Vect	= Vector
Dir	= Directory		
DMA	= Direct Memory Address		
Err	= Error		
FCB	= File Control Block		

Uppercase mnemonics refer to Data Structures; see the function definition. A . before a Data Structure means the byte offset of the Data Structure. A Return Code is either 0 for success or 0FFFFH to indicate failure. When the Return Code in AX is 0FFFFH, CX is the Error Code (see Table 6-5). An error code returned in AL is specific to the BDOS system call that was made.

Table 6-4. Data Structures Index

<i>Figure</i>	<i>Title</i>	<i>Page</i>
2-1	FCB - File Control Block	2-10
2-2	Directory Label Format	2-18
2-3	XFCB - Extended File Control Block	2-20
2-4	Directory Record with SFCB	2-24
2-5	SFCB Subfields	2-24
2-6	Disk System Reset	2-41
3-1	CMD File Header Format	3-3
3-2	Group Descriptor Format	3-3
3-3	Concurrent CP/M Base Page Values	3-6
4-1	Initial Program Stack	4-2
4-2	Concurrent CP/M 8080 Memory Model	4-3
4-3	Concurrent CP/M Small Memory Model	4-4
4-4	Concurrent CP/M Compact Memory Model	4-5
4-5	Intel Hexadecimal File Formats	4-10
5-1	8080 and Small RSP Models	5-2
5-2	RSP Header Format	5-3
5-3	RSP Command Queue Message	5-5
5-4	RSP Data Segment	5-7
6-1	ACB - Assign Control Block	6-21
6-2	Console Buffer Format	6-33
6-3	Drive, R/O, or Login Vector Structure	6-44
6-4	DPB - Disk Parameter Block	6-48
6-5	Disk Free Space Field Format	6-63
6-6	PFCB - Parse Filename Control Block	6-86
6-7	MCB - Memory Control Block	6-128
6-8	MPB - Memory Parameter Block	6-129
6-9	MFPB - M_FREE Parameter Block	6-132
6-10	APB - Abort Parameter Block	6-139

Table 6-4. (continued)

<i>Figure</i>	<i>Title</i>	<i>Page</i>
6-11	CLI Command Line Buffer	6-142
6-12	PD - Process Descriptor	6-146
6-13	UDA - User Data Area	6-151
6-14	CPB - Call Parameter Block	6-159
6-15	QPB - Queue Parameter Block	6-163
6-16	QD - Queue Descriptor	6-168
6-17	BDOS Version Number Format	6-174
6-18	BIOS Descriptor Format	6-175
6-19	Operating Systems Version Number Format	6-176
6-20	SERIAL Number Format	6-177
6-21	SYSDAT Table	6-179
6-22	TOD Time-of-Day Structure	6-185

Table 6-5. CX Error Code Reports

<i>Dec</i>	<i>Hex</i>	<i>Error Report</i>
0	00H	No error
1	01H	System call not implemented
2	02H	Illegal system call number
3	03H	Cannot find memory
4	04H	Illegal flag number
5	05H	Flag overrun
6	06H	Flag underrun
7	07H	No unused Queue Descriptors
8	08H	No free queue buffer
9	09H	Cannot find queue
10	0AH	Queue in use
12	0CH	No free process descriptors
13	0DH	No queue access
14	0EH	Empty queue
15	0FH	Full queue
16	10H	CLI queue missing
17	11H	No 8087 in system
18	12H	No unused Memory Descriptors
19	13H	Illegal console number

Table 6-5. (continued)

<i>Dec</i>	<i>Hex</i>	<i>Error Report</i>
20	14H	No Process Descriptor match
21	15H	No console match
22	16H	No CLI process
23	17H	Illegal disk number
24	18H	Illegal filename
25	19H	Illegal filetype
26	1AH	Character not ready
27	1BH	Illegal memory descriptor
28	1CH	Bad return from BDOS load
29	1DH	Bad return from BDOS read
30	1EH	Bad return from BDOS open
31	1FH	Null command
32	20H	Not owner of resource
33	21H	No CSEG in load file
34	22H	Process Descriptor exists on Thread Root
35	23H	Could not terminate process
36	24H	Cannot attach to process
37	25H	Illegal list device number
38	26H	Illegal password
40	28H	External termination occurred
41	29H	Fixup error upon load
42	2AH	Flag set ignored.

6.2 Concurrent CP/M System Calls

This section presents detailed information on the Concurrent CP/M system calls. Read the entire section through before attempting to use the system calls in a program, as many of them interact with one another.

6.2.1 Console I/O System Calls





Figure 6-1. ACB - Assign Control Block

Table 6-6. ACB Field Definitions

<i>Field</i>	<i>Definitions</i>
CNS	Console to assign
MATCH	Boolean; if OFFH, the process being assigned the console must have the CNS as its default console for a successful Assign. If 0H, no check is made.
PD	Process ID of the process being assigned the console. If this field is zero, a search is made of the Thread List for a process whose name is NAME. This field must be either zero or a valid Process ID. If this value is not a valid PD, an error occurs.
NAME	8-byte process name to search for. An error occurs if a process by this name does not exist.

The C_ASSIGN system call directly assigns the specified console to a specified process. This system call overrides the normal mechanism of the C_ATTACH and C_DETACH system calls. The system call returns an error code if a process other than the calling process owns the console. The system call ignores other processes waiting to attach to the specified console, and they continue to wait until the current owner either calls the C_DETACH system call, or terminates.

Refer to Table 6-5 for a list of error codes returned in CX.

C_ATTACH
Attach Default Console To Calling Process
Entry Parameters: Register CL: 092H (146)

The C_ATTACH system call attaches the default console to the calling process. If the console is already owned by the calling process or if it is not owned by another process, the C_ATTACH system call immediately returns with ownership established and verified. If another process owns the console, the calling process waits until the console becomes available.

Refer to Table 6-5 for a list of error codes returned in CX.

C_CATTACH	
Conditionally Attach Default Console To Calling Process	
Entry Parameters:	
Register CL:	0A2H (162)
Returned Values:	
Register AX:	0 if attach 'OK' OFFFH on failure
BX:	Same as AX
CX:	Error Code

The C_CATTACH system call attaches the default console of the calling process only if the console is currently unattached.

If the console is currently attached to another process, the system call returns a value of OFFH indicating that the console could not be attached. The system call returns a value of 0 to indicate that either the console is already attached to the process or that it was unattached and a successful attach operation was made.

Refer to Table 6-5 for a list of error codes returned in CX.

C_DELIMIT	
Set Or Return Output Delimiter	
Entry Parameters:	
Register CL:	06EH (110)
DX:	0FFFFH (get) or
DL:	Output Delimiter (set)
Returned Values:	
Register AL:	Output Delimiter or (no value if set)
BL:	Same as AL

A program can set or interrogate the current Output Delimiter by calling C_DELIMIT. If register DX = 0FFFFH, then the current Output Delimiter is returned in register AL. Otherwise, C_DELIMIT sets the Output Delimiter to the value in register DL.

C_DELIMIT sets the string delimiter for C_WRITESTR. When a new process is created, the default delimiter value is set to a dollar sign, \$. The default delimiter is not inherited from the parent process.

C_DETACH	
Detach Default Console From Calling Process	
Entry Parameters:	
Register CL:	093H (147)
Returned Values:	
Register AX:	0 if detach 'OK' 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The C_DETACH system call detaches the default console from the calling process. If the default console is not attached to the calling process, no action is taken. If other processes are waiting to attach to the console, the process with the highest priority attaches the console. If there is more than one process with the same priority waiting for the console, it is given to the queue writing processes on a first-come, first-serve basis.

Refer to Table 6-5 for a list of error codes returned in CX.

C_GET
Return The Calling Process's Default Console
Entry Parameters: Register CL: 099H (153)
Returned Values: Register AL: Console number BL: Same as AL

The C_GET system call returns the default console number of the calling process.

C_MODE	
Set Or Return Console mode	
Entry Parameters:	
Register CL:	06DH (109)
DX:	0FFFFH (get) or Console Mode (set)
Returned Values:	
Register AX:	Console Mode or (no value)
BX:	Same as AX

A process can set or interrogate the Console Mode by calling C_MODE. If register DX = 0FFFFH, then the current Console Mode is returned in register AX. Otherwise, C_MODE sets the Console Mode to the value contained in register DX.

The Console Mode is a 16-bit system parameter that determines the action of certain Console I/O functions. Note that the Console Mode bits are numbered from right to left. The Console Mode is set to zero when a new process created; it is not inherited from its parent. The definition of the Console Mode is

- bit 0 = 1 - CTRL-C only status for C_STAT.
= 0 - Normal status for C_STAT.

- bit 1 = 1 - Disable stop scroll, CTRL-S, start scroll, CTRL-Q, support.
= 0 - Enable stop scroll, start scroll support.

- bit 2 = 1 - Raw console output mode. Disables tab expansion for C_WRITE,
C_WRITESTR, and C_WRITEBLK. Also disables printer echo,
CTRL-P, support.
= 0 - Normal console output mode.

- bit 3 = 1 - Disable CTRL-C program termination
= 0 - Enable CTRL-C program termination

- bit 7 = 1 - Disable CTRL-O console output byte bucket
= 0 - Enable CTRL-O console output byte bucket

C_RAWIO	
Perform Direct Console I/O With Default Console	
Entry Parameters:	
Register CL:	06H (6)
DL:	0FFH (Input/ Status) or 0FEH (Status) or 0FDH (Input) or Character (Output)
Returned Values:	
Register AL:	(Input/Status) = 0H (No Character) = Character (Status) = 0H - No Character = 0FFH - Ready (Input) = Character (Output) No return value
BL:	Same as AL

The C_RAWIO system call allows the calling process to do raw console I/O to its default console. Concurrent CP/M verifies that the calling process owns its default console before allowing any I/O.

A process calls the C_RAWIO system call by passing one of three different values shown in Table 6-7.

Table 6-7. C_RAWIO Calling Values

<i>Value</i>	<i>Description</i>
0FFH	Console input status command (if no character is ready, a 00H is returned, else the character is returned).
0FEH	Console status command (on return, register AL contains 00H if no character is ready; otherwise it contains 0FFH).
0FDH	Console input command (if no character is ready, the calling process waits until one is typed). Input characters are not echoed to the screen.
ASCII character	If the parameter is less than 0FDH, C_RAWIO system call assumes register DL contains a valid ASCII character and sends it to the console.

The C_RAWIO system call places the calling process in Raw mode. The CTRL-C, CTRL-P, CTRL-S, and CTRL-O characters are not acted on by the PIN (Physical Input Process) but are passed on to the calling process when C_RAWIO is used.

Note: If the virtual console is in CTRL-S mode, and the process that owns the virtual console then performs a C_RAWIO call, the CTRL-S state is reset. Characters read with C_RAWIO are not echoed on the screen, thus allowing passwords and so forth to be entered in a secure manner.

C_READ	
Read A Character From The Default Console	
Entry Parameters:	
Register CL:	01H (1)
Returned Values:	
Register AL:	Character
BL:	Same as AL

The C_READ system call reads a character from the default console of the calling process. Before attempting the read, Concurrent CP/M internally verifies the ownership of the console. If the calling process does not own the console, it relinquishes the CPU resource until the calling process can attach to the console. Typically, a process that is created through the P_CLI system call owns its default console when it begins execution.

C_READ echoes characters read from the console. This includes the carriage return, line feed, and backspace characters. It expands tab characters (CTRL-I) in columns of eight characters.

C_READ ignores the termination character (CTRL-C) if the calling process cannot terminate (refer to the P_TERM system call). C_READ does not return until a character is typed on the console. The system suspends the calling process until a character is ready.

C_READSTR
Read An Edited Line From The Default Console
Entry Parameters: Register CL: 0AH (10) DX: BUFFER Address - Offset DS: BUFFER Address - Segment

The C_READSTR system call reads characters from the calling process's default console and places them into the specified buffer. The format of the buffer is shown in Figure 6-2. C_READSTR performs line-editing system calls on the line as it is read from the console; it completes a line and returns upon receiving a terminator character (carriage return or line feed) from the console or when the maximum number of characters is reached. As in the C_READ system call, C_READSTR echoes all graphic characters read from the console. Concurrent CP/M verifies that the calling process owns its default console before allowing I/O to begin.




Figure 6-2. Console Buffer Format

Table 6-8. Console Buffer Field Definition

<i>Field</i>	<i>Definition</i>
MAX	Maximum number of characters that can be read into the buffer. This value must be initialized before calling the C_READSTR system call.
NCHAR	Actual number of characters read into the buffer as filled in by the C_READSTR system call.
CHARACTERS	Actual characters read from the console as filled in by the C_READSTR system call.

C_READSTR recognizes a number of special characters used in editing the input line, as well as a set of special characters that actually control the calling process.

Table 6-9. C_READSTR Line-editing Characters

<i>Character</i>	<i>Function</i>
RUB/DEL	Removes the last character from the line and echoes it.
(CTRL-E)	Echoes new line, a carriage return (CTRL-M), and a line feed (CTRL-J), to the screen but does not affect the line buffer.
BACKSPACE (CTRL-H)	Removes the last character from the line and backspaces over that character.
TAB (CTRL-I)	Echoes enough spaces to place the next character position at a tab stop. Tab stops are fixed at every eighth character of the physical line.

Table 6-9. (continued)

<i>Character</i>	<i>Function</i>
LINE FEED (CTRL-J)	Terminates the input line. The C_READSTR system call does not echo a terminating character, nor does it place the character in the line buffer.
RETURN (CTRL-M)	Terminates the input line.
REDRAW (CTRL-R)	Retypes the current line after echoing a new line.
(CTRL-U)	Removes all of the current line from the line buffer, echoes a new line, and starts all over again.
(CTRL-X)	Removes all of the current line from the line buffer and echoes enough backspaces to return to the beginning of the line.

C_SET	
Set The Calling Process's Default Console	
Entry Parameters:	
Register CL:	094H (148)
DL:	Console Number
Returned Values:	
Register AX:	0 if successful 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The C_SET system call changes the calling process's default console to the value specified. If the console number specified is not one supported by this particular implementation of Concurrent CP/M, the system call returns an error code, and does not change the default console.

Refer to Table 6-5 for a list of error codes returned in CX.

C_STAT
Obtain The Status Of The Default Console
Entry Parameters: Register CL: 0BH (11)
Returned Values: Register AL: 01H character ready 00H not ready Register BL: Same as AL

The C_STAT system call checks to see if a character has been typed at the default console. If the calling process is not attached to its default console, the C_STAT system call causes a dispatch to occur and return 00H (the Not Ready condition).

This system call sets the console to the Nonraw mode, allowing recognition of special control characters such as the terminate character, CTRL-C. Use C_RAWIO to obtain console status in Raw mode.

Note: If bit 0 is set in the Console Mode word, using the C_MODE function call, C_STAT only returns AL = 01H when a CTRL-C is typed on the default console.

C_WRITE
Write A Character To The Default Console
Entry Parameters: Register CL: 02H (2) DL: ASCII character

The C_WRITE system call writes the specified character to the calling process's default console. As in the C_READ system call, Concurrent CP/M verifies that the calling process owns its default console before performing the operation. On output, C_WRITE expands tabs in columns of eight characters.

C_WRITEBLK
Send Specified String To Default Console:
Entry Parameters: Register CL: 06FH (111) DX: CHCB Address

C_WRITEBLK sends the character string located by the Character Control Block, CHCB, addressed in register pair DX to the console. If the Console Mode is in the Default state C_WRITEBLK expands tab characters, CTRL-I, in columns of eight characters.

The CHCB format is

- bytes 0 - 1 : Offset of character string
- bytes 2 - 3 : Segment of character string
- bytes 4 - 5 : Length of character string to print

C_WRITESTR
Print An ASCII String To The Default Console
Entry Parameters: Register CL: 09H (9) DX: STRING Address - Offset DS: STRING Address - Segment

The C_WRITESTR system call prints an ASCII string starting at the indicated string address and continuing until it reaches a dollar sign (\$) character (024H). \$ is the default string delimiter, and can be changed by the C_DELIMIT system call. C_WRITESTR writes this string to the calling process's default console.

Concurrent CP/M verifies that the calling process owns the console before writing the string. C_WRITESTR sets the console to a Nonraw state and expands tabs in columns of eight characters, as does the C_WRITE system call.

Use the C_WRITESTR system call whenever possible, rather than the single-character system calls. The CPU overhead involved in handling the first character is the same as that for a single-character system call, but subsequent characters require as little as one-fifth the CPU overhead.

6.2.2 Device System Calls

DEV_POLL	
Poll A Device	
Entry Parameters:	
Register CL:	083H (131)
DL:	Device Number
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The DEV_POLL system call is used by the XIOS to poll non interrupt-driven devices. It should be used whenever the XIOS is waiting for a non interrupt event. The calling process relinquishes the CPU and allows Concurrent CP/M to poll the device at every dispatch. The XIOS contains routines for each polling device number. These routines are called through the DEV_POLL system call, and they return whether the device is ready or not. When the device is ready, DEV_POLL restores the calling process to the RUN state and returns. Upon return, the calling process knows the device is ready.

Refer to Table 6-5 for a list of error codes returned in CX.

DEV_SETFLAG	
Set A System Flag	
Entry Parameters:	
Register CL:	085H (133)
DL:	Flag Number
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The DEV_SETFLAG system call is used by interrupt routines to notify the system that a logical interrupt has occurred. A process waiting for this flag is placed back into the RUN state. If there are no processes waiting, then the next process to wait for this flag returns successfully without relinquishing the CPU. The system call detects an error if the flag has already been set, and no process has done a DEV_WAITFLAG call to reset it.

Note: If a process waiting for a specific flag to be set is aborted, the next DEV_SETFLAG call is ignored and an error code is returned in CX. In this case, the interrupt handler should continue to set call DEV_SETFLAG until it successfully sets the flag IP, and AX = 0 on return.

Refer to Table 6-5 for a list of error codes returned in CX.

DEV_WAITFLAG	
Wait For A System Flag	
Entry Parameters:	Register CL: 084H (132) DL: Flag Number
Returned Values:	Register AX: 0 on success 0FFFFH on failure BX: Same as AX CX: Error Code

The DEV_WAITFLAG system call is used by a process to wait for an interrupt. The process relinquishes the CPU until an interrupt routine calls the DEV_SETFLAG system call, which places the waiting process in the RUN state. When DEV_WAITFLAG returns to the calling process, the interrupt has occurred, or an error has occurred. An error occurs when a process is already waiting for the flag. If the flag was set before DEV_WAITFLAG was called, the routine returns successfully without relinquishing the CPU. This routine is usually used by the XIOS. The mapping between types of interrupts and flag numbers is maintained in the XIOS, although Concurrent CP/M reserves flags 0, 1, 2, and 3 for system use.

Refer to Table 6-5 for a list of error codes returned in CX.

6.2.3 Disk Drive System Calls

The Drive Vector, Read-Only Vector, and Login Vectors are referenced or returned by several Concurrent CP/M Disk Drive system calls. The Drive, RO, or Login Vectors are 16-bit values specifying one or more drives, where the least significant bit corresponds to drive A, and the high-order bit corresponds to the sixteenth drive, labeled P. The format of the Drive, RO, and Login Vectors is illustrated below:




Figure 6-3. Drive, RO, or Login Vector Structure

DRV_ACCESS	
Access Specified Disk Drives	
Entry Parameters:	
Register CL:	026H (38)
DX:	Drive Vector
Returned Values:	
Register AL:	Return Code
AH:	Extended Error
BX:	Same as AX

The DRV_ACCESS system call inserts a special open file item into the system Lock List for each specified drive. While the item exists in the Lock List, the drive cannot be reset by another process. The calling process passes the drive vector in register DX. The format of the drive vector is discussed at the beginning of Section 6.2.3.

The DRV_ACCESS system call inserts no items if insufficient free space exists in the Lock List to support all the new items or if the number of items to be inserted puts the calling process over the Lock List open file maximum. If the BDOS Error mode is in the default mode (refer to the F_ERRMODE system call), the file system displays a message at the console identifying the error and terminates the calling process. Otherwise, DRV_ACCESS returns to the calling process with register AL set to 0FFH and register AH set to one of the following hexadecimal values.

0AH - Open File Limit Exceeded

0BH - No Room in system Lock List

On successful calls, DRV_ACCESS returns with register AL set to 00H.

DRV_ALLOCVEC
Get Allocation Vector Address For The Calling Process's Default Disk
Entry Parameters: Register CL: 01BH (27)
Returned Values: Register AX: ALLOC Address - Offset BX: Same as AX ES: ALLOC Address - Segment

Concurrent CP/M maintains an allocation vector in memory for each active disk drive. Some programs use the information provided by the allocation vector to determine the amount of free data space on a drive. Note, however, that the allocation information can be inaccurate if the drive has been marked Read-Only.

The DRV_ALLOCVEC system call returns the address of the allocation vector for the currently selected drive. If a physical error is encountered when the BDOS Error mode is in one of the return modes (refer to the F_ERRMODE system call), DRV_ALLOCVEC returns the value 0FFFFH in AX.

You can use the DRV_SPACE system call to directly return the number of free 128-byte records on a drive. The Concurrent CP/M utility, SHOW, finds a drive's free space by using the DRV_SPACE system call.

DRV_ALLRESET
Restore All Drives To Reset State
Entry Parameters: Register CL: 0DH (13)
Returned Values: Register AL: 0 if successful 0FFH on error Register BL: Same as AL

The DRV_ALLRESET system call restores the file system to a reset state where all the disk drives are set to Read-Write (refer to the DRV_SETRO and DRV_ROVEC system calls), the default disk is set to drive A, and the default DMA address is reset to offset 080H relative to the current DMA segment address. This system call can be used, for example, by an application program that requires disk changes during operation. You can also use the DRV_RESET system call for this purpose.

This system call is conditional under Concurrent CP/M. If another process has a file open on any of the drives to be reset, and the drive is also Read-Only or removable, the DRV_ALLRESET system call is denied, and none of the specified drives are reset (see Section 2.17).

Upon return, if the reset operation is successful, DRV_ALLRESET sets register AL to 00H. Otherwise, it sets register AL to 0FFH. If the BDOS is not in one of the return error modes (refer to the F_ERRMODE system call), the file system displays an error message at the console identifying the process owning the first open file that caused the DRV_ALLRESET to be denied.

DRV_DPB	
Return Address Of Disk Parameter Block For Calling Process's Default Disk	
Entry Parameters:	Register CL: 01FH (31)
Returned Values:	Register AX: DPB Address - Offset 0FFFFH on Physical Error
	BX: Same as AX
	ES: DPB Address - Segment

DRV_DPB returns the address of the XIOS-resident Disk Parameter Block (DPB) for the currently selected drive. The calling process can use this address to extract the disk parameter values.

If a physical error is encountered when the BDOS Error mode is one of the Return Error modes (refer to the F_ERRMODE system call), DRV_DPB returns the value 0FFFFH.

The Disk Parameter Block (DPB) contains the parameters that define the actual disk.




Figure 6-4. DPB - Disk Parameter Block

Table 6-10. DPB Field Definitions

<i>Field</i>	<i>Definition</i>
SPT Sectors Per Track	The number of Sectors Per Track equals the total number of physical sectors per track. Physical sector size is defined by PSH and PRM described below.
BSH Allocation Block Shift Factor	
BLM Allocation Block Mask	The data allocation block size determines the values of the data allocation Block Shift Factor and the allocation Block Mask. The Block Shift factor equals the logarithm base two of the block logical size in 128-byte records, or $BSH = \text{LOG}_2(\text{BLS})$. The Block Mask equals the number of 128-byte records in an allocation block minus 1, or $BLM = (2^{**}BSH) - 1$. Refer to the <i>Concurrent CP/M System Guide</i> for valid block sizes and BSH and BLM values.
EXM Extent Mask	The data block allocation size and the number of disk allocation blocks determine the value of the Extent Mask. The Extent Mask determines the maximum number of 16K extents that can be contained in a directory entry. It is equal to the maximum number of 16K extents per directory entry minus one. Refer to the <i>Concurrent CP/M System Guide</i> for EXM values.
DSM Disk Storage Maximum	The Disk Storage Maximum defines the total storage capacity of the drive. This is equal to the total number of allocation blocks minus 1 for the drive. DSM must be less than or equal to 7FFFH. If the disk uses 1024 byte blocks (BSH = 3, BLM = 7), DSM must be less than or equal to 00FFH.

Table 6-10. (continued)

<i>Field</i>	<i>Definition</i>
DRM Directory Maximum	The Directory Maximum defines the total number of directory entries for the drive. This is equal to the total number of directory entries, minus 1, that can be kept on this drive. The directory requires 32 bytes of disk per entry. The maximum directory allocation is 16 blocks, where the block size is determined by BSH and BLM.
AL0 Directory Allocation Vector 0 AL1 Directory Allocation Vector 1	The Directory Allocation Vectors determine the reserved directory allocation blocks.
CKS Checksum Vector Size	<p>The Checksum Vector Size determines the required length of the directory checksum vector and the number of directory entries that the BDOS will checksum. The Checksum Vector Size is equal to the number of directory entries divided by 4, or $CKS = (DRM + 1)/4$. If the media is fixed, CKS might be zero, no storage needs to be reserved, and the BDOS does not calculate directory checksums for the drive.</p> <p>The high-bit of CKS (that is, $\geq 08000H$) is set if the referenced drive is considered to be a nonremovable media drive. Note that this modifies the rules for resetting the drive. For more information, refer to Section 2.15.</p>

Table 6-10. (continued)

<i>Field</i>	<i>Definition</i>
OFF Track Offset	<p>The Track Offset is the number of reserved tracks at the beginning of the disk. OFF is equal to the track number on which the directory starts.</p>
PSH Physical Record Shift Factor	<p>The Physical Record Shift Factor ranges from 0 to 5, corresponding to physical record sizes of 128, 256, 512, 1K, 2K, or 4K bytes. It is equal to the logarithm base two of the physical record size divided by 128, or $\text{LOG}_2(\text{sector_size}/128)$.</p>
PRM Physical Record Mask	<p>The Physical Record Mask ranges from 0 to 31, corresponding to physical record sizes of 128, 256, 512, 1K, 2K, or 4K bytes. It is equal to the physical sector size divided by 128 minus 1, or $(\text{sector_size}/128) - 1$.</p> <p>For more information on DPB parameters, refer to the <i>Concurrent CP/M System Guide</i>, Section 5.4.</p>

DRV_FLUSH	
Flush Write-Deferred Buffers	
Entry Parameters:	
Register CL:	030H (48)
DL:	Purge Flag
Returned Values:	
Register AL:	Error Flag
AH:	Permanent Error
BX:	Same as AX

The DRV_FLUSH system call forces the write of any write-pending records contained in internal blocking/deblocking buffers. If register DL is set to 0FFH, DRV_FLUSH also purges all active data buffers after performing the writes. Programs that provide write with read verify support needed to purge internal buffers to ensure that verifying reads actually access the disk instead of returning data resident in internal data buffers. The Concurrent CP/M PIP utility is an example of such a program.

Upon return, the system call sets register AL to 00H if the flush operation is successful. If a physical error is encountered, DRV_FLUSH performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error and terminates the calling process. Otherwise, it returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical error codes:

01H - Disk I/O Error : permanent error

02H - Read/Only Disk

DRV_FREE
Free Specified Disk Drives
Entry Parameters: Register CL: 027H (39) DX: Drive Vector

The DRV_FREE system call purges the system Lock List of all file and locked record items that belong to the calling process on the specified drives. DRV_FREE passes the drive vector in register DX.

DRV_FREE does not close files associated with purged open file Lock List items. In addition, if a process references a purged file with a BDOS system call requiring an open FCB, the system call returns a checksum error. A file that has been written to should be closed before making a DRV_FREE call to the file's drive, or data can be lost. Refer to Section 2.17 for more information on this system call.

DRV_GET
Return The Calling Process's Default Drive
Entry Parameters: Register CL: 019H (25)
Returned Values: Register AL: Drive Number BL: Same as AL

The DRV_GET system call returns the calling process's currently selected default disk number. The disk numbers range from 0 through 15, corresponding to drives A through P.

DRV_GETLABEL	
Return Directory Label Data Byte For The Specified Drive	
Entry Parameters:	Register CL: 065H (101) DL: Drive
Returned Values:	Register AL: Directory Label Data Byte AH: Physical Error BX: Same as AX

The DRV_GETLABEL system call returns the directory label data byte for the specified drive. The calling process passes the drive number in register DL with 0 for drive A, 1 for drive B, continuing through 15 for drive P in a full 16-drive system. The format of the directory label data byte is shown below:

- bit 7 - Require passwords for password protected files
- 6 - Perform access time and date stamping
- 5 - Perform update time and date stamping
- 4 - Perform create time and date stamping
- 0 - Directory label exists on drive

(Bit 0 is the least significant bit)

DRV_GETLABEL returns the directory label data byte to the calling process in register AL. Register AL equal to 00H indicates that no directory label exists on the specified drive. If the system call encounters a physical error when the BDOS Error mode is in one of the return error modes (refer to the F_ERRMODE system call), it returns with register AL set to 0FFH and register AH set to one of the following:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error

DRV_LOGINVEC
Return Bit Map Of Logged-in Disk Drives
Entry Parameters: Register CL: 018H (24)
Returned Values: Register AX: Login Vector Register BX: Same as AX

The DRV_LOGINVEC system call returns the Login Vector in register AX. The Login Vector is a 16-bit value with the least significant bit corresponding to drive A, and the high-order bit corresponding to the 16th drive, drive P. A 0 bit indicates that the drive is not logged-in, while a 1 bit indicates the drive is logged in. Refer to the beginning of Section 6.2.3 for a complete description of the Login Vector.

DRV_RESET	
Reset Specified Disk Drives	
Entry Parameters:	
Register CL:	025H (37)
DX:	Drive Vector
Returned Values:	
Register AL:	Return Code
BL:	Same as AL

The DRV_RESET system call is used to programmatically restore specified removable media drives to the reset state (a reset drive is not logged in and is in Read-Write status). The passed parameter in register DX is a 16-bit vector of drives to be reset, where the least significant bit corresponds to drive A, and the high-order bit corresponds to the sixteenth drive, labeled P. Bit values of 1 indicate that the specified drive is to be reset. Refer to Section 2.17 for more information regarding the use of this system call.

This system call is conditional under Concurrent CP/M. If another process has a file open on any of the drives to be reset, the DRV_RESET system call is denied, and none of the drives are reset.

Upon return, if the reset operation is successful, DRV_RESET sets register AL to 00H. Otherwise, it sets register AH to 0FFH. If the BDOS Error mode is not in Return Error mode (refer to the F_ERRMODE system call), the system displays an error message at the console, identifying the process owning the first open file that caused the DRV_RESET request to be denied.

DRV_ROVEC
Return Bit Map Of Read-Only Disks
Entry Parameters: Register CL: 01DH (29)
Returned Values: Register AX: RO Vector BX: Same as AX

The DRV_ROVEC system call returns a bit vector indicating which drives have the temporary Read-Only bit set. The Read-Only bit can only be set by a DRV_SETRO call.

Note: When the file system detects a change in the media on a drive, it automatically logs in the drive and sets it to Read-Write.

The format of the RO Vector is analogous to that of the Login Vector. The least significant bit corresponds to drive A; the most significant bit corresponds to drive P. For a complete description of the RO Vector, refer to the beginning of this section.

DRV_SET	
Set Calling Process's Default Disk	
Entry Parameters:	
Register CL:	0EH (14)
DL:	Selected disk
Returned Values:	
Register AL:	Error Flag
AH:	Physical Error
BX:	Same as AX

The DRV_SET system call designates the specified disk drive as the default disk for subsequent BDOS file operations. Set the DL register to 0 for drive A, 1 for drive B, continuing through 15 for drive P. DRV_SET also logs in the designated drive if it is currently in the reset state. Logging in a drive activates the drive's directory for file operations.

FCBs that specify drive code zero (DR = 00H) automatically reference the currently selected default drive. FCBs with drive code values between 1 and 16, however, ignore the selected default drive and directly reference drives A through P.

Upon return, register AL equal to 00H indicates the select operation was successful. If a physical error is encountered, DRV_SET performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call).

If the BDOS Error mode is in the default mode, the system displays a message at the console, identifying the error and terminates the calling process. Otherwise, DRV_SET returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical error codes:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error

DRV_SETLABEL	
Create Or Update A Directory Label	
Entry Parameters:	Register CL: 064H (100) DX: FCB Address - Offset DS: FCB Address - Segment
Returned Values:	Register AL: Directory Code AH: Physical or Extended Error BX: Same as AX

The DRV_SETLABEL system call creates a directory label or updates the existing directory label for the specified drive. The calling process passes the address of an FCB containing the name, type, and extent fields to be assigned to the directory label. The name and type fields of the referenced FCB are not used to locate the directory label in the directory; they are simply copied into the updated or created directory label. Byte 12 of the FCB contains the user's specification of the directory label data byte.

The definition of the directory label data byte is

- bit 7 - Require passwords for password protected files
- 6 - Perform access time and date stamping
- 5 - Perform update time and date stamping
- 4 - Perform create time and date stamping
- 0 - Assign a new password to the directory label

(Bit 0 is the least significant bit)

If the current directory label is password protected, the correct password must be placed in the first 8 bytes of the current DMA or have been previously established as the default password (refer to the F_PASSWD system call). If bit 0 of the directory label data byte is set to 1, it indicates that a new password for the directory label has been placed in the second eight bytes of the current DMA.

The DRV_SETLABEL system call also requires that the referenced directory contains SFCBs in order to activate date and time stamping on the drive. If an attempt is made to activate date and time stamping when no SFCBs exist, the DRV_SETLABEL system call returns an error code and performs no action. The Concurrent CP/M INITDIR utility initializes a directory for date and time stamping by placing an SFCB in every fourth entry of the directory.

Upon return, the DRV_SETLABEL system call returns a directory code in register AL with the value 00H if the directory label create or update was successful, or 0FFH if no space existed in the referenced directory to create a directory label. It also returns 0FFH if date and time stamping was requested and the referenced directory did not contain SFCBs. Register AH is set to 00H in all of these cases.

If a physical or extended error is encountered, the DRV_SETLABEL system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the file system displays a message at the console identifying the error and terminates the calling process. Otherwise, the DRV_SETLABEL system call returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 04H - Invalid Drive : drive select error
- 07H - Password Error

DRV_SETRO
Set Default Disk To Read-Only
Entry Parameters: Register CL: 01CH (28)
Returned Values: Register AL: Return Code Register BL: Same as AL

The DRV_SETRO system call provides temporary write protection for the currently selected disk by marking the drive as Read-Only. No process can write to a disk that is in the Read-Only state. You must perform a successful DRV_RESET operation to restore a Read-Only drive to the Read-Write state (refer to the DRV_ALLRESET and DRV_RESET system calls).

The DRV_SETRO system call is conditional under Concurrent CP/M. If another process has an open file on the drive, the operation is denied, and the system call returns the value 0FFH to the calling process. Otherwise, it returns a 00H. If the BDOS Error mode is not in Return Error mode (refer to the F_ERRMODE system call), the file system displays an error message at the console, identifying the process owning the first open file that caused the DRV SETRO request to be denied.

Note that a drive in the Read-Only state cannot be reset by a process if another process has an open file on the drive.

DRV_SPACE	
Return Free Disk Space On Specified Drive	
<p>Entry Parameters:</p> <p style="padding-left: 40px;">Register CL: 02EH (46)</p> <p style="padding-left: 40px;">DL: Drive</p>	
<p>Returned Values:</p> <p style="padding-left: 40px;">Register AL: Error Flag</p> <p style="padding-left: 40px;">AH: Physical Error</p> <p style="padding-left: 40px;">BX: Same as AX</p> <p style="padding-left: 80px;">First 3 bytes of DMA Buffer filled in</p>	

The DRV_SPACE system call determines the number of free sectors (128-byte records) on the specified drive. The calling process passes the drive number in register DL, with 0 for drive A, 1 for B, continuing through 15 for drive P. DRV_SPACE returns a binary number in the first 3 bytes of the current DMA buffer. This number is returned in the format shown in Figure 6-5.



FS0 = LOW BYTE
 FS1 = MIDDLE BYTE
 FS2 = HIGH BYTE

Figure 6-5. Disk Free Space Field Format

Note that the returned free space value might be inaccurate if the drive has been marked Read-Only.

Upon return, DRV_SPACE sets register AL to 00H, indicating the operation was successful. However, if the BDOS Error mode is one of the return modes (refer to the F_ERRMODE system call), and a physical error occurs, it sets register AL to 0FFH, and register AH to one of the following values:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error

6.2.4 File-Access System Calls

Most file-access system calls reference a File Control Block (FCB). This data structure is illustrated in Table 2.1. Refer to Section 2.4 for a comprehensive explanation of the FCB data structure, its initialization, and usage.

F_ATTRIB	
Set The Attributes Of A Disk File	
Entry Parameters:	
Register CL:	01EH (30)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
BL:	Same as AL

By calling the F_ATTRIB system call, a process can modify a file's attributes and set its last record byte count. Other BDOS system calls can interrogate these file parameters, but only F_ATTRIB can change them. The file attributes that can be set or reset by F_ATTRIB are F1' through F4', Read-Only (T1'), System (T2'), and Archive (T3'). The specified FCB contains a filename with the appropriate attributes set or reset. The calling process must ensure that it does not specify an ambiguous filename. Also, if the specified file is password protected, the correct password must be placed in the first eight bytes of the current DMA buffer or have been previously established as the default password (refer to the F_PASSWD system call).

Interface attribute F5' specifies whether an extended file lock is to be maintained after the F_ATTRIB call. Interface attribute F6' specifies if the specified file's byte count is to be set. The interface attribute definitions are listed below:

- F5' = 0 - Do not maintain an extended file lock (default)
- F5' = 1 - Maintain an extended file lock
- F6' = 0 - Do not set byte count (default)
- F6' = 1 - Set byte count

If F5' is set and the referenced FCB specifies a file with an extended file lock, the calling process maintains the lock on the file. Otherwise, the file becomes available to other processes on the system. Section 2.11 describes extended file locking in detail.

If interface attribute F6' is set, the calling process must set the CR field of the referenced FCB to the new byte count value. A process can access a file's byte count value with the BDOS F_OPEN, F_SFIRST, and F_SNEXT system calls. File byte counts are described in section 2.15.

F_ATTRIB searches the FCB specified directory for an entry belonging to the current user number that matches the FCB specified name and type fields. The system call then updates the directory to contain the selected indicators, and if interface attribute F6' is set, the specified byte count value. Note that the last record byte count is maintained in the byte 13 of a file's directory FCBs.

File attributes T1', T2', and T3' are defined by Concurrent CP/M as described in Section 2.4.2. Attributes F1' through F4' of command files are defined as Compatibility Attributes, as described in Section 2.12. However, for all other files, attributes F1' through F4' are available for definition by the user. Attributes F5' through F8' are reserved as Interface Attributes and cannot be used as file attributes. Interface attributes are described in Section 2.4.3.

An F_ATTRIB system call is not performed if the referenced FCB specifies a file currently open for another process. It is performed, however, if the referenced file is open by the calling process in Locked mode. However, the file's lock entry is purged when this is done and the file system prevents continued read and write operations on the file. F_ATTRIB does not set the attributes of a file currently open in Read-Only or Unlocked mode for any process.

Making an F_ATTRIB system call for an open file can adversely affect the performance of the calling process. For this reason, you should close an open file before you call the F_ATTRIB system call.

Upon return, F_ATTRIB returns a directory code in register AL with the value 00H if the system call is successful, or 0FFH if the file specified by the referenced FCB is not found. Register AH is set to 00H in both cases.

If a physical or extended error is encountered, the F_ATTRIB system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the file system displays a message at the console identifying the error and terminates the process. Otherwise, it returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 04H - Invalid Drive : drive select error
- 05H - File open by another process
- 07H - Password Error
- 09H - Illegal ? in FCB

F_CLOSE	
Close A Disk File	
Entry Parameters:	
Register CL:	010H (16)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_CLOSE system call performs the inverse operation of the F_OPEN system call. The referenced FCB must have been previously activated by a successful F_OPEN or F_MAKE system call. Interface attributes F5' and F6' specify how the file is to be closed, as shown below:

- F5' = 0, F6' = 0 - Default Close
- F5' = 0, F6' = 1 - Extend File Lock
- F5' = 1, F6' = 0 - Partial Close
- F5' = 1, F6' = 1 - Partial Close

The F_CLOSE system call performs the following steps regardless of the interface attribute specification. First, it verifies that the referenced FCB has a valid checksum. If the checksum is invalid, F_CLOSE performs no action and returns an error code.

If the checksum is valid and the referenced FCB contains new information because of write operations to the FCB, F_CLOSE permanently records the new information in the directory. If the FCB does not contain new information, the directory update step is bypassed. However, F_CLOSE always attempts to locate the FCB's corresponding entry in the directory and returns an error code if the directory entry cannot be found.

If the F_CLOSE system call successfully performs the above steps, it performs different actions, depending on how the interface attributes are set. In default close operations, F_CLOSE decrements the file's open count, which is maintained in the file's system Lock List entry. If the open count decrements to zero, it indicates that the number of default close operations for the file matches the number of open operations.

If the open count decrements to zero, F_CLOSE permanently closes the file by performing the following steps. First of all, it removes the file's item from the system Lock List. If the FCB is opened in Unlocked mode, it also purges all record locks belonging to the file from the system Lock List. In addition, F_CLOSE invalidates the FCB's checksum to ensure the referenced FCB is not subsequently used with BDOS system calls that require an open FCB (for example, F_WRITE).

If the open count does not decrement to zero, F_CLOSE simply returns to the calling process and the file remains open.

For partial close operations, F_CLOSE does not decrement the file's open count and returns to the calling process. The file always remains open following a partial close request.

Closing a file with an extended file lock modifies the way F_CLOSE performs a permanent close. F_CLOSE only honors an extended lock request on a permanent close of a file opened in Locked mode. If these conditions are satisfied, F_CLOSE invalidates the FCB's checksum but maintains the lock item. Thus, although the file is permanently closed, other processes cannot access the file. Section 2.11 describes extended file locking in detail.

Upon return, the F_CLOSE system call returns a directory code in register AL with the value 00H if the close operation is successful, or 0FFH if the file is not found. Register AH is set to 0 in both of these cases.

If a physical or extended error is encountered, the F_CLOSE system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the file system displays a message identifying the error at the console and terminates the calling process. Otherwise the F_CLOSE system call returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 04H - Invalid Drive : drive select error
- 06H - Close Checksum Error

F_DELETE	
Delete A Disk File	
Entry Parameters:	
Register CL:	013H (19)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_DELETE system call removes files and/or XFCBs that match the FCB addressed in register DX. The filename and filetype fields can contain wildcard file specifications (question marks in bytes 1 through 11), but byte 0 cannot be a wildcard as it can be in the F_SFIRST and F_SNEXT system calls. Interface attribute F5' specifies the type of delete operation to be performed, as shown below:

F5' = 0 - Standard Delete (Default mode)

F5' = 1 - Delete only XFCB's and maintain an extended file lock.

If any of the files specified by the referenced FCB are password protected, the correct password must be placed in the first eight bytes of the current DMA buffer or it must have been previously established as the default password (refer to the F_PASSWD system call).

For standard delete operations, the F_DELETE system call removes all directory entries belonging to files that match the referenced FCB. All disk directory and data space owned by the deleted files is returned to free space and becomes available for allocation to other files. Directory XFCBs that were owned by the deleted files are also removed from the directory. If interface attribute F5' of the FCB is set to 1, F_DELETE deletes only the directory XFCBs matching the referenced FCB.

Note: If any of the files matching the input FCB specification fail the password check, are Read-Only, or are currently open by another process, then F_DELETE deletes no files or XFCBs. This applies to both types of delete operations.

Interface attribute F5' also specifies whether an extended file lock is to be maintained after the F_DELETE call. If F5' is set and the referenced FCB specifies a file with an extended lock, the calling process maintains the lock on the file. Section 2.11 describes extended file locking in detail.

A process can delete a file that it currently has open if the file is opened in locked mode. However, the BDOS returns a checksum error if the process makes a subsequent reference to the file with a BDOS system call requiring an open FCB. A process cannot delete files open in Read-Only or Unlocked mode.

Deleting an open file can adversely affect the performance of the calling process. For this reason, you should close an open file before you delete it.

Upon return, the F_DELETE system call returns a directory code in register AL with the value 00H if the delete is successful, or 0FFH if no file matching the referenced FCB is found. Register AH is set to 0 in both of these cases. If a physical or extended error is encountered, F_DELETE performs different actions, depending on the BDOS Error mode (refer to the F_ERRMODE system call).

If the BDOS Error mode is the default mode, the system displays a message identifying the error at the console and terminates the calling process. Otherwise, it returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 03H - Read-Only File
- 04H - Invalid Drive : drive select error
- 05H - File opened by another process or open in Read-Only or Unlocked mode
- 07H - Password Error

F_DMAGET
Return Address Of Direct Memory Access Buffer
Entry Parameters: Register CL: 034H (52)
Returned Values: Register AX: DMA Offset BX: Same as AX ES: DMA Segment

F_DMAGET returns the current DMA Base Segment address in ES, with the current DMA Offset in AX.

F_DMAOFF
Set The Direct Memory Address Offset
Entry Parameters: Register CL: 01AH (26) DX: DMA Address - Offset

DMA is an acronym for Direct Memory Address, which is often used with disk controllers that directly access the memory of the computer to transfer data to and from the disk subsystem. Under Concurrent CP/M, the current DMA is usually defined as the buffer in memory where a record resides before a disk write and after a disk read operation. If the BDOS Multisector Count is equal to one (refer to the F_MULTISEC system call), the size of the buffer is 128 bytes. However, if the BDOS Multisector Count is greater than one, the size of the buffer must equal $N * 128$, where N equals the Multisector Count.

Some BDOS system calls also use the current DMA to pass parameters and to return values. For example, BDOS system calls that check and assign file passwords require that the password be placed in the current DMA Buffer. As another example, DRV_SPACE returns its results in the first 3 bytes of the current DMA. When the current DMA is used in this context, the size of the buffer in memory is determined by the specific requirements of the system call.

When the P_CLT system call initiates a transient program, it sets the DMA offset to 080H and the DMA Segment or Base to its initial Data Segment. DRV_ALLRESET also sets the DMA offset to 080H. The F_DMAOFF system call can change this default value to another memory address. The DMA address remains at its current value until it is changed by an F_DMASEG, F_DMAOFF, or DRV_ALLRESET call.

F_DMASEG
Set Direct Memory Access Segment Address
Entry Parameters: Register CL: 033H (51) DX: DMA Segment Address

F_DMASEG sets the segment value of the current DMA buffer address. The word parameter in DX is a paragraph address and is used with the DMA offset value to specify the 20-bit address of the DMA buffer. Refer to the F_DMAOFF system call for additional information.

Note that upon initial program loading, the default DMA base is set to the address of the user's data segment (the initial value of DS) and the DMA offset is set to 080H, which provides access to the default buffer in the Base Page.

F_ERRMODE
Set BDOS Error Mode For Error Returns
Entry Parameters: Register CL: 02DH (45) DL: BDOS Error mode

The BDOS Error mode is a system parameter maintained for each running process that determines how the file system handles physical and extended errors. Physical and extended errors are described in Section 2.18. The BDOS Error mode has three states: the default mode, Return Error mode, and Return and Display mode.

If a physical or extended error occurs when the BDOS Error mode is in the default mode, the BDOS displays a system message at the console identifying the error and terminates the calling process.

If a physical or extended error occurs when the BDOS Error mode is in Return Error mode, the BDOS sets register AL to 0FFH, places an error code identifying the physical or extended error in register AH, and returns to the calling process.

If a physical or extended error occurs when the BDOS Error mode is in Return and Display mode, the BDOS displays the system message before returning to the calling process, and sets registers AH and AL as in the Return Error mode.

The F_ERRMODE system call sets the BDOS Error mode for the calling process to the mode specified in register DL. If register DL is set to 0FFH, the mode is set to Return Error mode. If register DL is set to 0FEH, the mode is set to Return and Display mode. If register DL is set to any other value, the mode is set to the default mode.

F_LOCK	
Lock Records In A Disk File	
Entry Parameters:	
Register CL:	02AH (42)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Error Code
AH:	Physical Error
BX:	Same as AX

The F_LOCK system call allows a process to establish temporary ownership to particular records within a file. This system call is only supported for files open in Unlocked mode. If it is called for a file open in Locked or Read-Only mode, no locking action is performed and a successful result is returned. This provides compatibility between Concurrent CP/M and CP/M-86.

The calling process passes the address of an FCB in which the random record field is filled with the Random Record Number of the first record to be locked. The number of records to be locked is determined by the BDOS Multisector Count (refer to the F_MULTISEC system call). The current DMA must also contain the 2-byte File ID returned by F_OPEN or F_MAKE when the referenced FCB was opened. Note that the File ID is only returned by the F_OPEN and F_MAKE system call when the Open mode is Unlocked.

Interface attribute F5' specifies the type of lock to perform. Interface attribute F6' specifies whether records have to exist in order to be locked. The F_LOCK interface attribute definitions are listed below:

- F5' = 0 - Exclusive lock (default)
- F5' = 1 - Shared lock
- F6' = 0 - Lock existing records only (default)
- F6' = 1 - Lock logical records.

These options are described in detail in Section 2.14.

F_LOCK verifies that a locking conflict with another process does not exist for each of the records to be locked. In addition, if F_LOCK is called with attribute F6' reset, it also verifies that each record number to be locked exists within the specified file. Both tests are made before any records are locked.

Most F_LOCK requests require a new entry in the BDOS system Lock List. If there is insufficient space in the system Lock List to satisfy the lock request, or if the process record lock limit is exceeded, then F_LOCK does not lock any records and returns an error code to the calling process.

Upon return, the F_LOCK system call sets register AL to 00H if the lock operation is successful. Otherwise, register AL contains one of the following error codes:

- 01H - Reading unwritten data
- 03H - Cannot close current extent
- 04H - Seek to unwritten extent
- 06H - Random Record Number out of range
- 08H - Record locked by another process
- 0AH - FCB Checksum Error
- 0BH - Unlocked file verification error
- 0CH - Process record lock limit exceeded
- 0DH - Invalid File ID
- 0EH - No Room in system Lock List
- 0FFH - Physical error; refer to register AH

The system call returns error code 01H when it accesses a data block that has not been previously written.

The system call returns error code 03H when it cannot close the current extent prior to moving to a new extent.

The system call returns error code 04H when it accesses an extent that has not been created.

The system call returns error code 06H when byte 35 (R2) of the referenced FCB is greater than 3.

The system call returns error code 08H if the specified record is locked by another process with an incompatible lock type.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0BH if the BDOS cannot locate the referenced FCB's directory entry when attempting to verify that the FCB contains current information.

The system call returns error code 0CH if performing the lock request would require that the process consume more than the maximum allowed number of system Lock List entries.

The system call returns error code 0DH when an invalid File ID is placed at the beginning of the current DMA.

The system call returns error code 0EH when the system Lock List is full and performing the lock request would require at least one new entry.

The system call returns error code 0FFH if a physical error is encountered, and the BDOS Error mode is either Return Error mode or Return and Display Error mode (refer to the F_ERRMODE system call). If the Error mode is in the default mode, the system displays a message at the console identifying the physical error and terminates the calling process. When the system call returns a physical error to the calling process, it is identified by register AH as shown below:

01H - Disk I/O Error : permanent error

04H - Invalid Drive : drive select error

F_MAKE	
Create A Disk File	
Entry Parameters:	
Register CL:	016H (22)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_MAKE system call creates a new directory entry for a file under the current user number. It also creates an XFCB for the file if the referenced drive has a directory label that enables password protection on the drive, and the calling process assigns a password to the file.

The calling process passes the address of the FCB with byte 0 of the FCB specifying the drive, bytes 1 through 11 specifying the filename and filetype, and byte 12 set to the extent number. Byte 12, the EX field, is usually set to 00H. Byte 32 of the FCB, the CR field, must be initialized to 00H, before or after the F_MAKE call, if the intent is to write sequentially from the beginning of the file.

Interface attribute F5' specifies the mode in which the file is to be opened. Interface attribute F6' specifies whether a password is to be assigned to the created file. The interface attributes are summarized below:

- F5' = 0 - Open in Locked mode (default)
- F5' = 1 - Open in Unlocked mode
- F6' = 0 - Do not assign password (default)
- F6' = 1 - Assign password to created file

When attribute F6' is set to 1, the calling process must place the password in the first 8 bytes of the current DMA buffer and set byte 9 of the DMA buffer to the password mode. Note that F_MAKE only interrogates attribute F6' if the referenced drive's directory label has enabled password support. The XFCB Password mode is summarized below:

XFCB Password Mode

- Bit 7 - Read mode
- Bit 6 - Write mode
- Bit 5 - Delete mode

The F_MAKE system call returns with an error code if the referenced FCB names a file that currently exists in the directory under the current user number. If there is any possibility of duplication, an F_DELETE call should precede the F_MAKE call.

If the make file operation is successful, it activates the referenced FCB for record operations (opens the FCB) and initializes both the directory entry and the referenced FCB to an empty file. It also computes a checksum and assigns it to the FCB. BDOS system calls that require an open FCB (for example, F_WRITE) verify that the FCB checksum is valid before performing their operation. If the file is opened in Unlocked mode, F_MAKE also sets bytes R0 and R1 in the FCB to a two-byte value called the File ID. The File ID is a required parameter for the BDOS Lock Record and Unlock Record system calls. Note that the F_MAKE system call initializes all file attributes to 0.

The BDOS file system also creates an open file item in the system Lock List to record a successful F_MAKE operation. While this item exists, no other process can delete, rename, truncate, or set the file attributes of this file.

A creation and/or update stamp is made for the created file if the referenced drive contains a directory label that enables creation and/or update time and date stamping and the FCB extent number is equal to 0.

F_MAKE also creates an XFCB for the created file if the referenced drive contains a directory label that enables password protection, interface attribute F6' of the FCB is 1, and the FCB is an extent zero FCB. In addition, F_MAKE also assigns the password and password mode placed in the first nine bytes of the DMA to the XFCB.

Upon return, the F_MAKE system call returns a directory code in register AL with the value 00H if the make operation is successful, or 0FFH if no directory space is available. Register AH is set to 00H in both cases.

If a physical or extended error is encountered, the F_MAKE system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error and terminates the calling process. Otherwise, it returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 04H - Invalid Drive : drive select error
- 08H - File Already Exists
- 09H - Illegal ? in FCB
- 0AH - Open File Limit Exceeded
- 0BH - No Room in system Lock List

F_MULTISEC	
Set BDOS Multisector Count	
Entry Parameters:	
Register CL:	02CH (44)
DL:	Number of Sectors
Returned Values:	
Register AL:	Return Code
BL:	Same as AL

The F_MULTISEC system call provides logical record blocking under Concurrent CP/M. It enables a process to read and write from 1 to 128 logical records of 128 bytes at a time during subsequent BDOS read and write system calls. It also specifies the number of 128-byte records to be locked or unlocked by the F_LOCK and F_UNLOCK system calls.

F_MULTISEC sets the Multisector Count value for the calling process to the value passed in register DL. Once set, the specified Multisector Count remains in effect until the calling process makes another F_MULTISEC system call and changes the value. Note that the P_CLI system call sets the Multisector Count to one when it initiates a transient process.

The Multisector Count affects BDOS error reporting for the BDOS read and write system calls. With the exception of physical errors, if an error occurs during these system calls and the Multisector Count is greater than one, the system returns the number of records successfully processed in register AH.

Upon return, the system call sets register AL to 00H if the specified value is in the range of 1 to 128. Otherwise, it sets register AL to 0FFH.

F_OPEN	
Open A Disk File	
Entry Parameters:	Register CL: 0FH (15) DX: FCB Address - Offset DS: FCB Address - Segment
Returned Values:	Register AL: Directory Code AH: Physical or Extended Error BX: Same as AX

The F_OPEN system call activates the FCB for a file that exists in the disk directory under the currently active user number or user zero. The calling process passes the address of the FCB, with byte 0 of the FCB specifying the drive, bytes 1 through 11 specifying the filename and filetype, and byte 12 specifying the extent. Byte 12 is usually set to zero.

Interface attributes F5' and F6' of the FCB specify the mode in which the file is to be opened, as shown below:

- F5' = 0, F6' = 0 - Open in Locked mode (Default mode)
- F5' = 1, F6' = 0 - Open in Unlocked mode
- F5' = 0 or 1, F6' = 1 - Open in Read-Only mode

If the file is password protected in Read mode, the correct password must be placed in the first eight bytes of the current DMA or have been previously established as the default password (refer to the F_PASSWD system call). If the current record field of the FCB, CR, is set to 0FFH, the F_OPEN system call returns the byte count of the last record of the file in the CR field. The last record byte count for a file can be set using the F_ATTRIB system call.

Note: The calling process must set the CR field of the FCB to 00H if the file is to be accessed sequentially from the first record.

The F_OPEN system call performs the following steps for files opened in locked or Read-Only mode. If the current user is nonzero and the file to be opened does not exist under the current user number, the F_OPEN system call searches user 0 for the file. If the file exists under user 0 and has the system attribute (T2') set, the file is opened under user 0. The Open mode is automatically set to Read-Only when this is done.

The F_OPEN system call also performs the following action for files opened in locked mode. If the file has the Read-Only attribute (T1') set, the Open mode is automatically set to Read-Only. Note that Read-Only mode implies the file can be concurrently accessed by other processes if they also open the file in Read-Only mode.

If the open operation is successful, F_OPEN activates the user's FCB for record operations as follows: F_OPEN copies the relevant directory information from the matching directory FCB into bytes D0 through D15 of the FCB. It also computes a checksum and assigns it to the FCB. All BDOS system calls that require an open FCB (for example, F_READ) verify that the FCB checksum is valid before performing their operation.

If the file is opened in Unlocked mode, the F_OPEN system call sets bytes R0 and R1 of the FCB to a two-byte value called the File ID. The File ID is a required parameter for the F_LOCK and F_UNLOCK system calls. If the Open mode is forced to Read-Only, F_OPEN sets interface attribute F8' to 1 in the user's FCB. In addition, the system call sets attribute F7' to 1 if the referenced file is password protected in Write mode and the correct password was not passed in the DMA or did not match the default password. The BDOS does not support write operations for an activated FCB if interface attribute F7' or F8' is set to 1.

The BDOS file system also creates an open file item in the system Lock List to record a successful open file operation. While this item exists, no other process can delete, rename, or modify the file's attributes. In addition, this item prevents other processes from opening the file if the file is opened in Locked mode. It also requires that other processes match the file's Open mode if the file is opened in Unlocked or Read-Only mode. This item remains in the system Lock List until the file is permanently closed or until the process that opened the file terminates.

When the open operation is successful, the F_OPEN system call also makes an access time and date stamp for the opened file when the following conditions are satisfied: the referenced drive has a directory label that requests access date and time stamping, the FCB extent field is equal to zero, and the referenced drive is Read-Write.

Upon return, F_OPEN returns a directory code in register AL with the value 00H if the open is successful, or 0FFH if the file is not found. Register AH is set to 0 in both of these cases. If a physical or extended error is encountered, the F_OPEN system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message identifying the error at the console and terminates the process. Otherwise, F_OPEN returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error
- 05H - File is open by another process or by the current process in an incompatible mode
- 07H - Password Error
- 09H - Illegal ? in FCB
- 0AH - Open File Limit Exceeded
- 0BH - No Room in system Lock List

F_PARSE	
Parse An ASCII String And Initialize An FCB	
Entry Parameters:	
Register CL:	098H (152)
DX:	PFCB Address - Offset
DS:	PFCB Address - Segment
Returned Values:	
Register AX:	0FFFFH if error 0 if end of filename string 0 if end of lineaddress of next item to parse
BX:	Same as AX
CX:	Error Code



Figure 6-6. PFCB-Parse Filename Control Block

Table 6-11. PFCB Field Definitions

<i>Field</i>	<i>Description</i>
FILENAME	Offset of an ASCII file specification to parse. The offset is relative to the same Data Segment as the PFCB.
FCBADR	Offset of a File Control Block to initialize. The offset is relative to the same Data Segment as the PFCB.

The F_PARSE system call parses an ASCII file specification (FILENAME) and prepares a File Control Block (FCB). The calling process passes the address of a data structure called the Parse Filename Control Block, (PFCB) in registers DX and DS. The PFCB contains the offset of the ASCII filename string followed by the offset of the target FCB.

F_PARSE assumes the file specification to be in the following form

{D;} FILENAME { .TYP } {;PASSWORD}

where those items enclosed in curly brackets are optional.

The F_PARSE system call parses the first file specification it finds in the input string. First of all, it eliminates leading blanks and tabs. F_PARSE then assumes the file specification ends on the first delimiter it encounters that is out of context with the specific field it is parsing. For instance, if it finds a colon (:), and it is not the second character of the file specification, the colon delimits the whole file specification.

The F_PARSE system call recognizes the following characters as delimiters:

- space
- tab
- return
- null
- ;
 (semicolon) - except before password field- =
 (equal)- <
 (less than)- >
 (greater than)- .
 (period) - except after filename and before filetype- :
 (colon) - except before filename and after drive- ,
 (comma)- |
 (vertical bar)- [
 (left square bracket)-]
 (right square bracket)

If the F_PARSE system call encounters a nongraphic character in the range 1 through 31 not listed above, it treats the character as an error.

The F_PARSE system call initializes the specified FCB as shown in Table 6-12.

Table 6-12. FCB Initialization

<i>Byte number</i>	<i>Explanation</i>
byte 0	The drive field is set to the specified drive. If the drive is not specified, the default value is used. 0 = default, 1 = A, 2 = B, etc.
byte 1-8	The name is set to the specified filename. All letters are converted to uppercase. If the name is not eight characters long, the remaining bytes in the filename field are padded with blanks. If the filename has an asterisk (*), all remaining bytes in the filename field are filled in with question marks (?). The system call returns an error if the filename is more than eight bytes long.
byte 9-11	The type is set to the specified filetype. If no type is specified, the type field is initialized to blanks. All letters are converted to uppercase. If the type is not three characters long, the remaining bytes in the filetype field are padded with blanks. If an asterisk is encountered, all remaining bytes are filled in with question marks. The system call returns an error if the type field is more than 3 bytes long.
byte 12-15	Filled in with zeros.
byte 16-23	The password field is set to the specified password. If no password is specified, this field is initialized to blanks. If the password is not eight characters long, remaining bytes are padded with blanks. All letters are converted to uppercase. The system call returns an error if the password field is more than eight bytes long.
byte 24-31	Reserved for system use.

If an error occurs, F_PARSE returns 0FFFFH in register AX indicating the error.

On a successful parse, the F_PARSE system call checks the next item in the FILENAME string. It scans for the first character that follows trailing blanks and tabs. If the character is a line feed (0AH), a carriage return (0DH), or a null character (00H), it returns a 0 indicating the end of the FILENAME string. If the next character is a delimiter, it returns the address of the delimiter. If the next character is not a delimiter, it returns the address of the first trailing blank or tab.

If the F_PARSE system call is to be used to parse a subsequent filename in the FILENAME string, the returned address should be advanced over the delimiter before placing it in the PFCB.

Refer to Table 6-5 for a list of error codes returned in CX.

F_PASSWD
Establish A Default Password For File Access
Entry Parameters: Register CL: 06AH (106) DX: Password Address - Offset DS: Password Address - Segment

The F_PASSWD system call allows a process to specify a password value before a file protected by the password is accessed. When the file system accesses a password-protected file, it checks the current DMA, and the default password for the correct value. If either value matches the file's password, full access to the file is allowed.

Concurrent CP/M maintains a default password for each process running on the system. A new process inherits its initial default password from its parent, the process creating the new process.

Note: Changing the default password does not affect other processes currently running on the system.

To make an F_PASSWD call, the calling process passes the address of an eight-byte field containing the password.

F_RANDREC	
Return The Random Record Number Of The Next Record To Access In A Disk File	
Entry Parameters:	
Register CL:	024H (36)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	Random Record Field of FCB Set

The F_RANDREC system call returns the Random Record Number of the next record to be accessed from a file that has been read or written sequentially to a particular point. The system call returns this value in the Random Record field, bytes R0, R1, and R2, of the addressed FCB. The F_RANDREC system call can be useful in two ways.

First, it is often necessary to initially read and scan a sequential file to extract the positions of various key fields. As each key is encountered, F_RANDREC is called to compute the random record position for the data corresponding to this key. If the data unit size is 128 bytes, the resulting record number minus one is placed into a table with the key for later retrieval.

After scanning the entire file and tabularizing the keys and their record numbers, you can move directly to a particular record by performing a random read using the corresponding Random Record Number that was saved earlier. The scheme is easily generalized when variable record lengths are involved, because the program need only store the buffer-relative byte position along with the key and record number in order to find the exact starting position of the keyed data at a later time.

F_RANDREC can also be used when switching from a sequential read or write to a random read or write. A file is sequentially accessed to a particular point in the file, F_RANDREC is called to set the record number, and subsequent random read and write operations continue from the next record in the file.

F_READ	
Read Records Sequentially From A Disk File	
Entry Parameters:	
Register CL:	014H (20)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Error Code
AH:	Physical Error
BX:	Same as AX

The F_READ system call reads the next 1 to 128 128-byte records from a file into memory, beginning at the current DMA address. The BDOS Multisector Count (refer to the F_MULTISEC system call) determines the number of records to be read. The default is one record. The addressed FCB must have been previously activated by an F_OPEN or F_MAKE system call.

F_READ reads each record from the current record (CR) field in the FCB, relative to the current extent, then automatically increments the CR field to the next record position. If the CR field overflows, then F_READ automatically opens the next logical extent and resets the CR field to zero for the next read operation. The calling process must set the CR field to 00H following the open call if the intent is to read sequentially from the beginning of the file.

Upon return, the F_READ system call sets register AL to zero if the read operation is successful. Otherwise, register AL contains an error code identifying the error as shown below:

- 01H - Reading unwritten data (end-of-file)
- 08H - Record locked by another process
- 09H - Invalid FCB
- 0AH - FCB Checksum Error
- 0BH - Unlocked file verification error
- 0FFH - Physical error; refer to register AH

The system call returns error code 01H if no data exists at the next record position of the file. The no data situation is usually encountered at the end of a file. However, it can also occur if you try to read a data block that has not been previously written or an extent that has not been created. These situations are usually restricted to files created or appended with the BDOS random write system calls (F_WRITERAND and F_WRITEZF).

The system call returns error code 08H if the calling process attempts to read a record locked by another process with an exclusive lock. This error code is only returned for files opened in Unlocked mode.

The system call returns error code 09H if the FCB is invalidated by a previous F_CLOSE system call that returned an error.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0BH if the BDOS cannot locate the FCB's directory entry when attempting to verify that the referenced FCB contains current information. The system call only returns this error for files opened in Unlocked mode.

The system call returns error code 0FFH if a physical error is encountered and the BDOS Error mode is in one of the return modes (refer to the F_ERRMODE system call). If the Error mode is in the default mode, the file system displays a message at the console identifying the physical error and terminates the calling process. When the system call returns a physical error to the calling process, it is identified by register AH as shown below:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error

On all error returns, except for physical error returns (AL = 255), F_READ sets register AH to the number of records successfully read before the error was encountered. This value can range from 0 to 127 depending on the current BDOS Multisector Count. It is always set to zero when the Multisector Count is equal to one.

F_READRAND	
Read Random Records From A Disk File	
Entry Parameters:	Register CL: 021H (33)
	DX: FCB Address - Offset
	DS: FCB Address - Segment
Returned Values:	Register AL: Error Code
	AH: Physical Error
	BX: Same as AX

The F_READRAND system call is similar to the F_READ system call except that the read operation takes place at a particular Random Record Number, selected by the 24-bit value constructed from the three-byte, R0, R1, R2, field beginning at position 33 of the FCB. Note that the sequence of 24 bits is stored with the least significant byte first, R0, the middle byte next, R1, and the high byte last, R2. The Random Record Number can range from 0 to 262,143. This corresponds to a maximum value of 3 in byte R2.

To read a file with the F_READRAND system call, the calling process must first open the base extent, extent 0. This ensures that the FCB is properly initialized for subsequent random access operations. The base extent might or might not contain any allocated data.

The F_READRAND system call reads the record specified by the random record field into the current DMA address. F_READRAND automatically sets the FCB extent and current record number values, EX and CR, but unlike the F_READ system call, it does not advance the current record number. Thus, a subsequent F_READRAND call rereads the same record. After a random read operation, a file can be accessed sequentially, starting from the current randomly accessed position. However, the last randomly accessed record is reread or rewritten when switching from random to sequential mode.

If the BDOS Multisector count is greater than one (refer to the F_MULTISEC system call), F_READRAND reads multiple consecutive records into memory beginning at the current DMA. F_READRAND automatically increments the R0, R1, R2 field of the FCB to read each record. However, it restores the FCB's Random Record Number to the first record's value upon return to the calling process.

Upon return, F_READRAND sets register AL to 00H if the read operation is successful. Otherwise, register AL contains one of the following error codes:

- 01H - Reading unwritten data
- 03H - Cannot close current extent
- 04H - Seek to unwritten extent
- 06H - Random Record Number out of range
- 08H - Record locked by another process
- 0AH - FCB Checksum Error
- 0BH - Unlocked file verification error
- 0FFH - Physical error; refer to register AH

The system call returns error code 01H when it accesses a data block not previously written. This may indicate an end-of-file (EOF) condition.

The system call returns error code 03H when it cannot close the current extent prior to moving to a new extent.

The system call returns error code 04H when a read random operation accesses an extent that has not been created.

The system call returns error code 06H when byte 35 (R2) of the referenced FCB is greater than 3.

The system call returns error code 08H if the calling process attempts to read a record locked by another process with an exclusive lock. This error code is only returned for files opened in Unlocked mode.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0BH if the BDOS cannot locate the FCB's directory entry when attempting to verify that the referenced FCB contains current information. The system call only returns this error for files open in Unlocked mode.

The system call returns error code 0FFH if a physical error is encountered and the BDOS Error mode is in one of the return modes (refer to the F_ERRMODE system call). If the Error mode is in the default mode, the file system displays a message at the console identifying the physical error and terminates the calling process. When a physical error is returned to the calling process, it is identified by the four low-order bits of register AH as shown below:

01H - Disk I/O Error : permanent error

04H - Invalid Drive : drive select error

On all error returns except for physical error returns, AL = 255, F_READRAND sets register AH to the number of records successfully read before the error was encountered. This value can range from 0 to 127 depending on the current BDOS Multisector Count. It is always set to zero when the Multisector Count is equal to one.

F_RENAME	
Rename A Disk File	
Entry Parameters:	
Register CL:	017H (23)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_RENAME system call uses the referenced FCB to change all directory entries of the file specified by the drive and filename in bytes 0 to 11 of the FCB to the filename specified in bytes 17 through 27.

If the file specified by the first filename is password-protected, the correct password must be placed in the first eight bytes of the current DMA buffer, or have been previously established as the default password (refer to the F_PASSWD system call).

The calling process must also ensure that the filenames specified in the FCB are valid and unambiguous, and that the new filename does not already exist on the drive. F_RENAME uses the drive code at byte 0 of the FCB to select the drive. The drive code at byte 16 of the FCB is ignored.

Interface attribute F5' specifies whether an extended file lock is to be maintained after the F_ATTRIB call as shown below:

F5' = 0 - Do not maintain an extended file lock (default)

F5' = 1 - Maintain an extended file lock

If F5' is set and the referenced FCB specifies a file with an extended file lock, the calling process maintains the lock on the file. Otherwise, the file becomes available to other processes on the system. Section 2.11 describes extended file locking in detail.

A process can rename a file that it has open if the file is open in locked mode. However, the BDOS returns a checksum error if the process subsequently references the file with a system call requiring an open FCB. A file open in Read-Only or Unlocked mode cannot be renamed by any process.

Renaming an open file can adversely affect the performance of the calling process. For this reason, you should close an open file before you rename it.

Upon return, the F_RENAME system call returns a directory code in register AL with the value 00H if the rename is successful, or 0FFH if the file named by the first filename in the FCB is not found. Register AH is set to 00H in both of these cases. If a physical or extended error is encountered, the F_RENAME system call performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error, and terminates the process. Otherwise, it returns to the calling process with register AL set to 0FFH and with register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read-Only Disk
- 03H - Read-Only File
- 04H - Invalid Drive : drive select error
- 05H - File open by another process
- 07H - Password Error
- 08H - File Already Exists
- 09H - Illegal ? in FCB

F_SFIRST	
Find The First File That Matches The Specified FCB	
Entry Parameters:	
Register CL:	011H (17)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_SFIRST system call scans the directory for a match with the referenced FCB. Two types of searches can be performed. For standard searches, the calling process initializes bytes 0 through 12 of the referenced FCB, with byte 0 specifying the drive directory to be searched, bytes 1 through 11 specifying the file or files to be searched for, and byte 12 specifying the extent. Byte 12 is usually set to 00H. An ASCII question mark (63, or 03FH hexadecimal) in any of the bytes 1 through 12 matches all entries on the directory in the corresponding position. This facility, called ambiguous file reference, can be used to search for multiple files on the directory. When called in the standard mode, F_SFIRST scans for the first file entry in the specified directory that matches the FCB and belongs to the current user number.

The F_SFIRST system call also initializes the F_SNEXT system call. After the F_SFIRST system call has located the first directory entry matching the referenced FCB, F_SNEXT can be called repeatedly to locate all remaining matching entries. In terms of execution sequence, however, the F_SNEXT call must follow either a F_SFIRST or F_SNEXT call with no other intervening BDOS file-access system calls.

If byte 0 of the referenced FCB is set to a question mark, F_SFIRST ignores the remainder of the referenced FCB and locates the first directory entry residing on the current default drive. All remaining directory entries can be located by making multiple F_SNEXT calls. This type of search operation is not usually made by application programs, but it does provide complete flexibility to scan all directory entries. Note that this type of search operation must be performed to access a drive's directory label.

Upon return, the F_SFIRST system call returns a directory code in register AL with the value 0 to 3 if the search is successful, or 0FFH if a matching directory entry is not found. Register AH is set to zero in both of these cases. For successful searches, the current DMA is also filled with the directory record containing the matching entry, and the relative starting position is $AL * 32$. The directory information can be extracted from the buffer at this position.

If the directory has been initialized for date and time stamping, then an FCB resides in every fourth directory entry, and successful directory codes are restricted to the values 0 to 2. For successful searches, if the matching directory record is an extent zero entry, and if an SFCB resides at offset 96 within the current DMA buffer, then the contents of $(DMA\ Address + 96) = 021H$, and the SFCB contains the time and date stamp information and password mode for the file. This information is located at the relative starting position of $97 + (AL * 10)$ within the current DMA in the following format:

0 - 3 : Create or Access Date and Time Stamp Field
4 - 7 : Update Date and Time Stamp Field
8 : Password Mode Field

Refer to Section 2.8 for more information about SFCBs.

If a physical error is encountered, the F_SFIRST system call performs different actions depending on the BDOS error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message identifying the error at the console and terminates the calling process. Otherwise, it returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical error codes:

01H - Disk I/O Error : permanent error
04H - Invalid Drive : drive select error

F_SIZE	
Compute The Size Of A Disk File	
Entry Parameters:	Register CL: 023H (35) DX: FCB Address - Offset DS: FCB Address - Segment
Returned Values:	Register AL: Directory Code AH: Physical or Extended Error BX: Same as AX Random Record Field of FCB Set

The F_SIZE system call determines the virtual file size. This is the address of the record immediately following the end of the file. The virtual size of a file corresponds to the physical size if the file is written sequentially. If the file is written in random mode, gaps might exist in the allocation, and the file might contain fewer records than the indicated size. For example, if a single record with record number 262,143, the Concurrent CP/M maximum, is written to a file using the F_WRITERAND system call, then the virtual size of the file is 262,144 records even though only one data block is actually allocated.

To compute file size, the calling process passes the address of an FCB with bytes R0, R1, and R2 present. The F_SIZE system call sets the random record field of the FCB to the Random Record Number + 1 of the last record in the file. If the R2 byte is set to 04H, and R0 and R1 are both zero, then the file contains the maximum record count, 262,144.

A process can append data to the end of an existing file by calling F_SIZE to set the random record position to the end of file, and then performing a sequence of random writes.

Note: The file need not be open in order to use F_SIZE. However, if the file is open in Locked mode and it has been extended by the calling process, the file must be closed before F_SIZE is called. Otherwise, F_SIZE returns an incorrect file size. F_SIZE returns the correct size for files open in Unlocked mode and Read-Only mode.

Upon return, F_SIZE returns a 00H in register AL if the file specified by the referenced FCB is found, or a 0FFH in register AL if the file is not found. Register AH is set to 00H in both cases.

If a physical or extended error is encountered, F_SIZE performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error and terminates the process. Otherwise, F_SIZE returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error
- 09H - Illegal ? in FCB

F_SNEXT
Find A Subsequent File That Matches The Specified FCB Of A Previous F_SFIRST Or F_SNEXT
Entry Parameters: Register CL: 012H (18)
Returned Values: Register AL: Directory Code AH: Physical or Extended Error BX: Same as AX

The F_SNEXT system call is identical to F_SFIRST except that the directory scan continues from the last entry that was matched. F_SNEXT returns a directory code in register AL, analogous to F_SFIRST.

Note: In execution sequence, a F_SNEXT call must follow either an F_SFIRST or another F_SNEXT with no other intervening BDOS file-access system calls.

F_TIMEDATE
Return File Date Stamps And Password Mode
<p>Entry Parameters:</p> <p style="padding-left: 40px;">Register CL: 066H (102)</p> <p style="padding-left: 80px;">DX: FCB Address - Offset</p> <p style="padding-left: 80px;">DS: FCB Address - Segment</p> <p>Returned Values:</p> <p style="padding-left: 40px;">Register AL: Directory Code</p> <p style="padding-left: 80px;">AH: Physical Error</p> <p style="padding-left: 80px;">BX: Same as AX</p>

The F_TIMEDATE system call returns the time and date stamp information and password mode for the specified file in byte 12 and bytes 24 through 31 of the specified FCB. The calling process passes the address of an FCB in which the drive, filename, and type fields have been defined.

If F_TIMEDATE is successful, it sets the following fields in the referenced FCB

byte 12 password mode field

- bit 7 - Read mode
- bit 6 - Write mode
- bit 5 - Delete mode

Byte 12 equal to 0 indicates the file has not been assigned a password.

- byte 24 - 27 XFCB Create or Access time stamp field
- byte 28 - 31 XFCB Update time stamp field

Upon return, F_TIMEDATE returns a directory code in register AL with the value 00H if the operation is successful, or 0FFH if the specified file is not found. Register AH is set to 00H in both of these cases. If a physical or extended error is encountered, F_TIMEDATE performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error and terminates the calling process. Otherwise, F_TIMEDATE returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical error codes:

- 01H - Disk I/O Error : permanent error
- 04H - Invalid Drive : drive select error
- 09H - Illegal ? in FCB

F_TRUNCATE	
Truncate File	
Entry Parameters:	
Register CL:	063H (99)
DX:	FCB Address - Offset
Returned Values:	
Register AL:	Directory Code
AH:	Physical or Extended Error
BX:	Same as AX

The F_TRUNCATE system call sets the last record of a file to the Random Record Number contained in the referenced FCB. The calling program passes the address of the FCB in register DX with byte 0 of the FCB specifying the drive, bytes 1 through 11 specifying the filename and filetype, and bytes 33 through 35 (R0, R1, and R2) specifying the last record of the file. The last record number is a 24-bit value, stored with the least significant byte first (R0), the middle byte next (R1), and the high byte last (R2). This value can range from 0 to 262,143 (03FFFFH).

If the file specified by the referenced FCB is password-protected, the correct password must have been placed in the first eight bytes of the current DMA buffer, or have been previously established as the default password (refer to the F_PASSWD system call).

Interface attribute F5' specifies whether an extended file lock is to be maintained after the F_TRUNCATE call, as shown below:

F5' = 0 - Do not maintain an extended file lock (default)

F5' = 1 - Maintain an extended file lock

If F5' is set and the referenced FCB specifies a file with an extended file lock, the calling process maintains the lock on the file. Otherwise, the file becomes available to other processes on the system. Section 2.11 describes extended file locking in detail.

F_TRUNCATE requires that the Random Record Number field of the referenced FCB specify a value less than the current file size. In addition, if the file is sparse, the random record field must specify a region of the file where data exists.

A process can truncate a file that it currently has open if the file is opened in locked mode, and the file has not been extended during the open session. However, the BDOS returns a checksum error if the process makes a subsequent reference to the file with a BDOS system call requiring an open FCB. A process cannot truncate files open in RO or Unlocked mode.

Truncating an open file is not recommended under Concurrent CP/M. F_TRUNCATE truncates a file based on the file's state in the directory. If a process attempts to truncate a region of the file that has been allocated in memory but has not been recorded in the directory, F_TRUNCATE returns an error. Even when successful, an open file truncate can adversely affect the performance of the calling process. For these reasons, you should close an open file before you truncate it.

After completion, F_TRUNCATE returns a directory code in register AL with the value 00H if the operation is successful or 0FFH if the file is not found or if the record number is invalid. In both cases register AH is set to 00H.

If a physical or extended error is encountered, F_TRUNCATE performs different actions depending on the BDOS error mode (refer to F_ERRMODE). If the BDOS error mode is in the default mode, a message identifying the error is displayed at the console and the program is terminated. Otherwise, F_TRUNCATE returns to the calling program with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read/Only Disk
- 03H - Read/Only File
- 04H - Invalid Drive : drive select error
- 05H - File Currently Open
- 06H - Close Checksum Error
- 07H - Password Error
- 08H - File Already Exists
- 09H - Illegal ? in FCB
- 0AH - Open File Limit Exceeded
- 0BH - No Room in System Lock List

F_UNLOCK	
Unlock Records In A Disk File	
Entry Parameters:	
Register CL:	02BH (43)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Error Code
AH:	Physical Error
BX:	Same as AX

The F_UNLOCK system call unlocks one or more consecutive records previously locked by the F_LOCK system call. This system call is only supported for files open in Unlocked mode. If it is called for a file open in Locked or Read-Only mode, no unlocking action occurs and a successful result is returned. Record locking and unlocking is described in detail in Section 2.14.

The calling process passes the address of an FCB in which the Random Record Field is filled with the Random Record Number of the first record to be unlocked. The number of records to be unlocked is determined by the BDOS Multisector Count (refer to the F_MULTISEC system call). The current DMA must contain the 2-byte File ID returned by the F_OPEN or F_MAKE system call when the referenced FCB was opened. Note that the File ID is only returned by F_OPEN or F_MAKE when the file open mode is Unlocked.

If interface attribute F5' is set to 1, F_UNLOCK unlocks all locked records belonging to the calling process. The F_UNLOCK interface attribute definition is listed below:

- F5' = 0 - Unlock records specified by Random Record Number and BDOS Multisector Count (default)
- F5' = 1 - Unlock all locked records.

F_UNLOCK ignores the FCB Random Record field and the BDOS Multisector Count when F5' is set.

F_UNLOCK does not unlock a record that is currently locked by another process. However, the system call does not return an error if a process attempts to do that. Thus, if the Multisector Count is greater than one, F_UNLOCK unlocks all records locked by the calling process, skipping those records locked by other processes.

Some F_UNLOCK requests require a new entry in the BDOS system Lock List. If there is insufficient space in the system Lock List to satisfy the F_UNLOCK request, or if the process record Lock List limit is exceeded, then F_UNLOCK does not unlock any records and returns an error code to the calling process.

Upon return, F_UNLOCK sets register AL to 00H if the unlock operation was successful. Otherwise, register AL contains one of the following error codes:

- 01H - Reading unwritten data
- 03H - Cannot close current extent
- 04H - Seek to unwritten extent
- 06H - Random Record Number out of range
- 0AH - FCB Checksum Error
- 0CH - Process record Lock List limit exceeded
- 0DH - Invalid File ID
- 0EH - No room in system Lock List
- 0FFH - Physical error refer to register AH

The system call returns error code 01H when it accesses a data block which has not been previously written.

The system call returns error code 03H when it cannot close the current extent prior to moving to a new extent.

The system call returns error code 04H when it accesses an extent that has not been created.

The system call returns error code 06H when byte 35 (r2) for a list of the referenced FCB is greater than 3.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0CH if performing the unlock request would require that the process consume more than the maximum allowed number of system Lock List entries.

The system call returns error code 0DH when an invalid File ID is placed at the beginning of the current DMA.

The system call returns error code 0EH when the system Lock List is full and performing the unlock request would require at least one new entry.

The system call returns error code 0FFH if a physical error was encountered and the BDOS Error mode is one of the return modes (refer to the F_ERRMODE system call). If the Error mode is the Default mode, the system displays a message at the console identifying the physical error and terminates the calling process. When the system call returns a physical error to the calling process, it is identified by register AH as shown below:

01H - Disk I/O Error : permanent error

04H - Invalid Drive : drive select error

F_USERNUM	
Set Or Return The Calling Process's Default User Number	
Entry Parameters:	Register CL: 020H (32) DL: 0FFH to GET User Number User Number to SET
Returned Values:	Register AL: Current User Number if GET BL: Same as AL

A process can change or interrogate its current default user number by calling F_USERNUM. If register DL = 0FFH, then the system call returns the value of this user number in register AL. The value can range from 0 to 0FH. If register DL is not 0FFH, then the system call changes the default user number to the value in DL, modulo 010H (the high nibble of DL is masked off).

Under Concurrent CP/M, a new process inherits its initial default user number from its parent, the process creating the new process. Changing the default user number does not change the user code of the parent. On the other hand, all child processes of the calling process inherit the new user number.

This convention is demonstrated by the operation of the TMP. When a command is typed, a new process is created with the same user number as that of the TMP. If this new process changes its user number, the TMP is unaffected. Once the new process terminates, the TMP displays the same user number in its prompt that it displayed before the command was entered and the child process was created.

F_WRITE	
Write Records Sequentially To A Disk File	
Entry Parameters:	
Register CL:	015H (21)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Error Code
AH:	Physical Error
BX:	Same as AX

The F_WRITE system call writes 1 to 128, 128-byte data records beginning at the current DMA address into the file named by the specified FCB. The BDOS Multisector Count (refer to the F_MULTISEC system call) determines the number of 128-byte records that are written. The default is one record. An F_OPEN or F_MAKE system call must have previously activated the referenced FCB.

F_WRITE places the record into the file at the position indicated by the CR byte of the FCB, and then automatically increments the CR byte to the next record position. If the CR field overflows, the system call automatically opens or creates the next logical extent and resets the CR field to 00H in preparation for the next write operation. If F_WRITE is used to write to an existing file, then the newly written records overlay those already existing in the file. The calling process must set the CR field to 00H following an F_OPEN or F_MAKE system call if the intent is to write sequentially from the beginning of the file.

F_WRITE makes an update date and time stamp for the file if the following conditions are met: the referenced drive has a directory label that requests update date and time stamping, and the file has not already been stamped for update by a previous F_MAKE or F_WRITE system call.

Upon return, the F_WRITE system call sets register AL to 00H if the write operation is successful. Otherwise, register AL contains an error code identifying the error as shown below:

- 01H - No available directory space
- 02H - No available data block
- 08H - Record locked by another process
- 09H - Invalid FCB
- 0AH - FCB Checksum Error
- 0BH - Unlocked file verification error
- 0FFH - Physical error; refer to register AH

The system call returns error code 01H when it attempts to create a new extent that requires a new directory entry, and no available directory entries exist on the selected disk drive.

The system call returns error code 02H when it attempts to allocate a new data block to the file, and no unallocated data blocks exist on the selected disk drive.

The system call returns error code 08H if the calling process attempts to write to a record locked by another process, or a record locked by the calling process in shared mode. The system call returns this error only for files open in Unlocked mode.

The system call returns error code 09H if the FCB is invalidated by a previous F_CLOSE system call that returned an error.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0BH if the BDOS cannot locate the FCB's directory entry when attempting to verify that the referenced FCB contains current information. The system call returns this error only for files open in Unlocked mode.

The system call returns error code 0FFH if a physical error was encountered and the BDOS is in Return Error mode or Return and Display Error mode (refer to the F_ERRMODE system call). If the Error mode is the Default mode, the system displays a message at the console identifying the physical error and terminates the calling process. When the system call returns a physical error to the calling process, it is identified by register AH as shown below:

- 01H - Disk I/O Error : permanent error
- 02H - Read/Only Disk
- 03H - Read/Only File or
File Opened in Read/Only Mode or
File password protected in Write mode
- 04H - Invalid Drive : drive select error

On all error returns except for physical error returns ($AL = 255$), F_WRITE sets register AH to the number of records successfully written before the error was encountered. This value can range from 0 to 127, depending on the current BDOS Multisector Count. It is always set to zero when the Multisector Count is equal to one.

F_WRITERAND	
Write Random Records To A Disk File	
Entry Parameters:	Register CL: 022H (34) DX: FCB Address - Offset DS: FCB Address - Segment
Returned Values:	Register AL: Error Code AH: Physical Error BX: Same as AX

The F_WRITERAND system call is analogous to the F_READRAND system call, except that data is written to the disk from the current DMA address. If the disk extent and/or data block where the data is to be written is not already allocated, the BDOS automatically performs the allocation before the write operation continues.

In order to write to a file using the F_WRITERAND system call, the calling process must first open the base extent, extent 0. This ensures that the FCB is properly initialized for subsequent random access operations. If the file is empty, the calling process must create the base extent with the F_MAKE system call before an F_WRITERAND system call. The base extent might or might not contain data, but it records the file in the directory so that it can be displayed by the DIR utility. If a process does not open extent 0 and allocates data to some other extent, the file is invisible to the DIR utility.

The F_WRITERAND system call sets the logical extent and current record positions to correspond with the random record being written, but does not change the Random Record Number. Thus sequential read or write operations can follow a random write, with the current record being reread or rewritten as the calling process switches from random to sequential mode.

F_WRITERAND makes an update date and time stamp for the file if the following conditions are met: the referenced drive has a directory label that requests update date and time stamping, and the file has not already been stamped for update by a previous F_MAKE or F_WRITE system call.

If the BDOS Multisector Count is greater than one (refer to the F_MULTISEC system call), the F_WRITERAND system call writes multiple consecutive records from memory beginning at the current DMA address. The system call automatically increments the R0, R1, and R2 field of the FCB to write each record. However, it restores the FCB's Random Record Number to the first record's value upon return to the calling process.

Upon return, the F_WRITERAND system call sets register AL to 00H if the write operation is successful. Otherwise, register AL contains one of the following error codes:

- 02H - No available data block
- 03H - Cannot close current extent
- 05H - No available directory space
- 06H - Random record number out of range
- 08H - Record locked by another process
- 0AH - FCB Checksum Error
- 0BH - Unlocked file verification error
- 0FFH - Physical error refer to register AH

The system call returns error code 02H when it attempts to allocate a new data block to the file. No unallocated data blocks exist on the selected disk drive.

The system call returns error code 03H when it cannot close the current extent before moving to a new extent.

The system call returns error code 05H when it attempts to create a new extent that requires a new directory entry and no available directory entries exist on the selected disk drive.

The system call returns error code 06H when byte 35 (R2) of the referenced FCB is greater than 3.

The system call returns error code 08H if the calling process attempts to write to a record locked by another process, or a record locked by the calling process in shared mode. The system call returns this error only for files open in Unlocked mode.

The system call returns error code 0AH if the referenced FCB failed the FCB checksum test.

The system call returns error code 0BH if the BDOS cannot locate the FCB's directory entry when attempting to verify that the referenced FCB contains current information. The system call returns this error only for files open in Unlocked mode.

The system call returns error code 0FFH if a physical error is encountered and the BDOS Error mode is in one of the return modes (refer to the F_ERRMODE system call). If the Error mode is in the default mode, the system displays a message at the console identifying the physical error and terminates the calling process. When a physical error is returned to the calling process, it is identified by register AH as shown below:

- 01H - Disk I/O Error : permanent error
- 02H - Read/Only Disk
- 03H - Read/Only File or
File Opened in Read/Only Mode or
File password protected in Write mode
- 04H - Invalid Drive : drive select error

On all error returns, except for physical error returns (AL = 255), F_WRITERAND sets register AH to the number of records successfully written before the error was encountered. This value can range from 0 to 127 depending on the current BDOS Multisector Count. It is always set to zero when the Multisector Count is equal to one.

F_WRITEXFCB	
Write Extended File Control Block Of A Disk File	
Entry Parameters:	Register CL: 067H (103) DX: FCB Address - Offset DS: FCB Address - Segment
Returned Values:	Register AL: Directory Code AH: Physical or Extended Error BX: Same as AX

The F_WRITEXFCB system call creates a new XFCB or updates the existing XFCB for the specified file. The calling process passes the address of an FCB in which the drive, name, type, and extent fields have been defined. The FCB extent field, if set, specifies the password mode and whether a new password is to be assigned to the file. The format of the extent field byte is shown below:

FCB byte 12 (EX) XFCB password mode

bit 7 - Read mode

bit 6 - Write mode

bit 5 - Delete mode

bit 0 - assign new password to the file

If the FCB is currently password-protected, the correct password must reside in the first 8 bytes of the current DMA or have been previously established as the default password (refer to the F_PASSWD system call). If bit 0 is set to 1, the new password must reside in the second 8 bytes of the current DMA.

Note: The F_WRITEXFCB system call does not create or update an XFCB if the XFCB specifies a file open by another process. However, a process can update or create an XFCB for a file that it has open in Locked mode.

Upon return, F_WRITEFCB returns a directory code in register AL with the value 00H if the XFCB create or update was successful. F_WRITEFCB returns 0FFH in register AL if no directory label existed on the specified drive, or the file specified in the FCB was not found, or no space existed in the directory to create an XFCB, or if the drive is not password enabled. F_WRITEFCB also returns 0FFH if passwords are not enabled by the specified drive's directory label. Register AH is set to 00H in all of these cases.

If a physical or extended error is encountered, F_WRITEFCB performs different actions depending on the BDOS Error mode (refer to the F_ERRMODE system call). If the BDOS Error mode is in the default mode, the system displays a message at the console identifying the error and terminates the calling process. Otherwise, F_WRITEFCB returns to the calling process with register AL set to 0FFH and register AH set to one of the following physical or extended error codes:

- 01H - Disk I/O Error : permanent error
- 02H - Read/Only Disk
- 04H - Invalid Drive : drive select error
- 05H - File open by another process, or open in Read-Only or Unlocked mode
- 07H - Password Error
- 09H - Illegal ? in FCB

F_WRITEZF	
Write A Random Record To A Disk File And Prefill New Data Blocks With Zeros	
Entry Parameters:	
Register CL:	028H (40)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AL:	Error Code
AH:	Physical Error
BX:	Same as AX

The F_WRITEZF system call is similar to the F_WRITERAND system call, with the exception that it fills a previously unallocated data block with zeros (00H) before writing the record. If this system call has been used to create a file, records accessed by an F_READRAND system call that contain all zeros identify unwritten random records. Unwritten random records in allocated data blocks of files created using the F_WRITERAND system call contain uninitialized data.

6.2.5 List Device I/O System Calls

L_ATTACH
Attach The Default List Device To The Calling Process
Entry Parameters: Register CL: 09EH (158)

The L_ATTACH system call attaches the default list device of the calling process. If the list device is already attached to some other process, the calling process relinquishes the CPU until the other process detaches from the list device. When the list device becomes free, and the calling process is the highest priority process waiting for the list device, the attach operation occurs.

Refer to Table 6-5 for a list of error codes returned in CX.

L_CATTACH
Conditionally Attach To The Default List Device
Entry Parameters: Register CL: 0A1H (161)
Returned Values: Register AX: 0 if attach 'OK' 0FFFFH on failure Register BX: Same as AX Register CX: Error Code

The L_CATTACH system call attaches the default list device of the calling process only if the list device is currently available.

If the list device is currently attached to another process, the system call returns a value of 0FFH, indicating that the list device could not be attached. The system call returns a value of 00H to indicate that either the list device is already attached to the process, or that it was unattached, and a successful attach operation was made.

Refer to Table 6-5 for a list of error codes returned in CX.

L_DETACH	
Detach The Default List Device From The Calling Process	
Entry Parameters:	
Register CL:	09FH (159)
Returned Values:	
Register AX:	0 if detach 'OK' 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The L_DETACH system call detaches the default list device of the calling process. If the list device is not currently attached, no action takes place.

Refer to Table 6-5 for a list of error codes returned in CX.

L_GET
Return The Calling Process's Default List Device
Entry Parameters: Register CL: 0A4H (164) Returned Values: Register AL: List Device Number BL: Same as AL

The L_GET system call returns the default list device number of the calling process.

L_SET
Set The Calling Process's Default List Device
Entry Parameters: Register CL: 0A0H (160) DL: List Device Number Returned Values: Register CX: Error Code

The L_SET system call sets the default list device for the calling process.

Refer to Table 6-5 for a list of error codes returned in CX.

L_WRITE
Write A Character To The Default List Device
Entry Parameters: Register CL: 05H (5) DL: Character

The L_WRITE system call writes the specified character to the default list device of the calling process. Before writing the character, the system internally calls L_ATTACH to verify that the calling process owns its default list device.

L_WRITEBLK
Send Specified Character String to Default List Device
Entry Parameters: Register CL: 070H (112) DX: CHCB Address

L_WRITEBLK sends the character string specified in the Character Control Block (CHCB) and addressed in register pair DX to the logical list device, LST:. The CHCB format is

- bytes 0 - 1 : Offset of character string
- bytes 2 - 3 : Segment of character string
- bytes 4 - 5 : Length of character string to print

6.2.6 Memory System Calls

There are two classes of Memory System Calls in Concurrent CP/M. The first class supports the MP/M-86 memory allocation scheme and contains two system calls, M_ALLOC and M_FREE. The second class contains six system calls, MC_ABS, MC_ALLFREE, MC_ALLOC, MC_ALLOCABS, MC_FREE, and MC_MAX. These system calls support the CP/M-86 memory allocation scheme.

Note: The CP/M-86 memory calls are also supported under MP/M-86.

Many of the Memory system calls use the Memory Control Block (MCB) or the Memory Parameter Block (MPB) to pass parameters to and from the operating system. The format, structure and example programming equates for these data structures are presented below, along with example listings.




Figure 6-7. MCB - Memory Control Block

Table 6-13. MCB Field Definitions

<i>Field</i>	<i>Definition</i>
BASE	The Segment Address of the beginning of the specified memory segment.
LENGTH	Length of the Memory Segment in paragraphs. The LENGTH field is set to the number of paragraphs wanted.
EXT	The EXT field is unused but must be available.

```

;*****
;
;*
;*      Memory Control Block Definition
;*
;*****
mcb_base      equ      word ptr 0
mcb_length    equ      word ptr mcb_base + word
mcb_ext       equ      byte ptr mcb_length + word

mcb_len       equ      mcb_ext + byte
;
    
```

Listing 6-1. Memory Control Block Definition




Figure 6-8. MPB - Memory Parameter Block

Table 6-14. MPB Field Definitions

<i>Field</i>	<i>Description</i>
START	if non-00H, an absolute request at this paragraph
MIN	minimum memory needed (paragraphs)
MAX	maximum memory wanted (paragraphs)
* 0000H	these fields must be 00H; they are used internally.

```

;*****
;*
;*      Memory Parameter Block Definition
;*
;*****
mpb_start      equ      word ptr 0
mpb_min        equ      word ptr mpb_start + word
mpb_max        equ      word ptr mpb_min + word
mpb_pdadr      equ      word ptr mpb_max + word
mpb_flags      equ      word ptr mpb_pdadr + word

mpb_len        equ      mpb_flags + word

; mpb_flags definition

mf_load        equ      00001h
mf_share       equ      00002h
mf_code        equ      00004h

;

```

Listing 6-2. Memory Parameter Block Definition

M_ALLOC	
Allocate A Memory Segment	
Entry Parameters:	
Register CL:	080H or 081H (128,129)
DX:	MPB Address-Offset
DS:	MPB Address-Segment MPB filled in
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code MPB_start filled in

The M_ALLOC system call allows a program to allocate extra memory. A successful allocation allocates a contiguous memory segment whose length is at least the MIN and no more than the MAX number of paragraphs specified in the MPB. The START field of the MPB is modified to be the starting paragraph of the memory segment. The MIN and MAX fields are modified to be the length of the memory segment in paragraphs. Memory Segments can be explicitly released through the M_FREE system call; Concurrent CP/M also releases all memory owned by a process at termination.

Note: MIN and MAX fields must be explicitly filled in. The MAX value must be greater than or equal to the MIN value.

Refer to Table 6-5 for a list of error codes returned in CX.

M_FREE	
Free A Memory Segment	
Entry Parameters:	
Register CL:	082H (130)
DX:	MFPB Address - Offset
DS:	MFPB Address - Segment
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code




Figure 6-9. MFPB - M_FREE Parameter Block

The M_FREE system call releases memory starting at the START paragraph to the end of a single previously allocated segment that contains the START paragraph. If the START paragraph is the same as that returned in the MPB of a memory allocation call, then M_FREE releases the whole memory segment. The * 0000H field must be initialized to zero.

Refer to Table 6-5 for a list of error codes returned in CX.

MC_ABSALLOC	
Allocate A Memory Segment At A Specified Address	
Entry Parameters:	
Register CL:	038H (56)
DX:	MCB Address - Offset
DS:	MCB Address - Segment
Returned Values:	
Register AL:	0 on success 0FFH on failure
BL:	Same as AL
CX:	Error Code

The MC_ABSALLOC system call allocates a memory area that starts at the address specified by the BASE field. The memory area's length is specified by the LENGTH field of the MCB. Upon return, register AL contains a 00H if the request was successful, and a 0FFH if the memory could not be allocated. If the calling process already owns the requested memory, no error is returned. This assures compatibility with CP/M-86.

Refer to Table 6-5 for a list of error codes returned in CX.

MC_ABSMAX	
Allocate Maximum Memory Available At A Specified Address	
Entry Parameters:	Register CL: 036H (54) DX: MCB Address - Offset DS: MCB Address - Segment MCB_base filled in, MCB_length set to max number of paragraphs wanted
Returned Values:	Register AL: 0 on success 0FFH on failure BL: Same as AL CX: Error Code MCB_length set to actual number of paragraphs allocated

In CP/M-86, system call 036H does not allocate memory, but under Concurrent CP/M, this system call allocates memory, because other processes are competing for common memory. For compatibility with CP/M-86, MC_ABSALLOC (system call 56) does not return an error if there is a memory segment allocated at the absolute address.

MC_ABSMAX is used to allocate the largest possible region at the absolute paragraph boundary given by the BASE field of the MCB, for a maximum of LENGTH paragraphs. If the allocation is successful, the system call sets the LENGTH to the actual length. Upon return, register AL has the value 0FFH if no memory is available at the absolute address, and 00H if the request was successful.

Refer to Table 6-5 for a list of error codes returned in CX.

MC_ALLFREE
Free All Memory Owned By The Calling Process
Entry Parameters: Register CL: 03AH (58)

In the Concurrent CP/M environment, the MC_ALLFREE system call releases all of the calling process's memory except the User Data Area (UDA). This system call is useful for system processes and for subprocesses that share the memory of another process.

Note: This system call should not be used by processes running programs loaded into the Transient Program Areas (TPAs).

MC_ALLOC	
Allocate A Memory Segment	
Entry Parameters:	
Register CL:	037H (55)
DX:	MCB Address - Offset
DS:	MCB Address - Segment
	MCB_length filled in
Returned Values:	
Register AL:	0 on success
	0FFH on failure
BL:	Same as AL
CX:	Error Code
	MCB_base filled in

The MC_ALLOC system call allocates a memory area whose size is the LENGTH field of the MCB. MC_ALLOC returns the base paragraph address of the allocated region in the user's MCB. Upon return, register AL contains a 00H if the request was successful, and a 0FFH if the memory could not be allocated.

Refer to Table 6-5 for a list of error codes returned in CX.

MC_FREE	
Free A Specified Memory Segment	
<p>Entry Parameters:</p> <p>Register CL: 039H (57)</p> <p>DX: MCB Address - Offset</p> <p>DS: MCB Address - Segment</p> <p style="padding-left: 40px;">MCB_base, MCB_ext filled in</p>	
<p>Returned Values:</p> <p>Register AL: 0 if successful</p> <p style="padding-left: 40px;">OFFH on failure</p> <p>BL: Same as AL</p> <p>CX: Error Code</p>	

The MC_FREE system call is used to release memory areas allocated to the program. The value of the EXT field of the MCB controls the operation of this system call. If EXT = OFFH, then the system call releases all memory areas allocated by the calling program. If the EXT field is 00H, the system call releases the memory area beginning at the specified BASE and ending at the end of the previously allocated memory segment.

Refer to Table 6-5 for a list of error codes returned in CX.

MC_MAX	
Allocate Maximum Memory Available	
Entry Parameters: Register CL: 035H (53) DX: MCB Address - Offset DS: MCB Address - Segment (MCB_length contains maximum number of paragraphs wanted)	
Returned Values: Register AL: 0 on success OFFH on failure BL: Same as AL CX: Error Code (MCB_base filled in, MCB_length set to actual number of paragraphs allocated)	

In CP/M-86, system call 035H does not allocate memory, but under Concurrent CP/M, this system call allocates memory because other processes are competing for common memory. For compatibility with CP/M-86, MC_ABSALLOC (system call 56) does not return an error if there is a memory segment allocated at the absolute address.

MC_MAX allocates the largest available memory region that is less than or equal to the LENGTH field of the MCB in paragraphs. If the allocation is successful, the system call sets the BASE to the base paragraph address of the available area and LENGTH to the paragraph length. Upon return, register AL has the value OFFH if no memory is available, and 00H if the request was successful. The system call sets the EXT to 1 if there is additional memory for allocation, and 0 if no additional memory is available.

Refer to Table 6-5 for a list of error codes returned in CX.

6.2.7 Process/Program System Calls

P_ABORT	
Terminate A Process By Name Or PD Address	
Entry Parameters: <ul style="list-style-type: none"> Register CL: 09DH (157) DX: APB Address - Offset DS: APB Address - Segment APB filled in 	
Returned Values: <ul style="list-style-type: none"> Register AX: 0 on success 0FFH on failure BX: Same as AX CX: Error Code 	




Figure 6-10. APB - Abort Parameter Block

Table 6-15. APB Field Definitions

<i>Field</i>	<i>Definition</i>
PD	Process Descriptor offset of the process to be terminated. If this field is zero, a match is attempted with the NAME and CNS fields to find the process. If this field is nonzero, the NAME and CNS fields are ignored.
TERM	Termination Code. This field corresponds to the termination code of the P_TERM system call. If the low-order byte of TERM is OFFH, P_ABORT can abort a specified system process; if the termination code is not OFFH, the system call can only terminate a user process. (A system process is identified by the SYS flag in the Process Descriptor's FLAG field.)
*00H	This field is reserved for system use and must be set to zero.
CNS	Default console of process to be aborted. If the PD field is 0, the P_ABORT system call scans the Thread List for a PD with the same NAME and CNS fields as specified in the APB. P_ABORT only aborts the first process that it finds. Subsequent calls must be made to abort all processes with the same NAME and CNS.
NAME	Name of the process to be aborted. Combined with the CNS field, the NAME field is used to find the process to be aborted. This is only used if the PD field is 0.

The P_ABORT system call permits a process to terminate another specified process. The calling process passes the address of a data structure called an Abort Parameter Block, initialized as described above.

If the Process Descriptor address is known, it can be filled in, and the process name and console can be omitted. Otherwise, the Process Descriptor address field should be a 00H and the process name and console must be specified. In either case, the calling process must supply the termination code, which is the same parameter passed to the P_TERM system call.

Refer to Table 6-5 for a list of error codes returned in CX.

P_CLI
Interpret And Execute Command Line
<p>Entry Parameters:</p> <p style="margin-left: 20px;">Register CL: 096H (150)</p> <p style="margin-left: 40px;">DX: CLBUF Address - Offset</p> <p style="margin-left: 40px;">DS: CLBUF Address - Segment</p> <p>Returned Values:</p> <p style="margin-left: 20px;">Register AX: 0 on success 0FFFFH on error</p> <p style="margin-left: 40px;">CX: Error Code</p>




Figure 6-11. CLI Command Line Buffer

Table 6-16. Command Line Buffer Field Definitions

<i>Field</i>	<i>Definition</i>
*00H	Must be set to zero for system use.
COMMAND	1-128 ASCII characters terminated with a null character.

The P_CLI system call obtains an ASCII command from the Command Line Buffer (CLBUF) and then executes it. If the calling process is attached to its default virtual console, the P_CLI system call assigns the virtual console to either the newly created process, or to the Resident System Process (RSP) that acts on the command. The calling process must reattach to its default virtual console before accessing it.

P_CLI calls F_PARSE to parse the command line. If an error occurs in F_PARSE, P_CLI returns to the calling process with the error code set to the same code that F_PARSE returned.

If there is no disk specification for the command, P_CLI tries to open a system queue with the same name as the command. If the open operation is successful, and the queue is an RSP-type queue, P_CLI then writes the command tail to the RSP queue. If the queue is full, the system call returns an error code to the calling process. The P_CLI function also attempts to assign the calling process's virtual console to a process with the same name as the RSP queue. If the RSP queue cannot be found, the CLI assumes the command is on disk and continues.

The P_CLI system call opens a file with the filename being the command and the filetype being CMD. If the command has an explicit disk specification, and the F_OPEN system call fails, P_CLI returns an error code to the calling process. If there is no disk specification with the command, P_CLI attempts to open the command file on the system disk. If the F_OPEN system call succeeds, P_CLI checks the file to verify the SYSTEM attribute is on. This search order is discussed in Section 2.9.1 of the *Concurrent CP/M User's Guide*. If this second F_OPEN fails or if the DIR attribute is on, P_CLI returns an error code to the calling process.

Once the P_CLI system call succeeds in opening the command file, it calls the P_LOAD system call. The P_LOAD system call finds, and then loads the file into an appropriate memory space. If P_LOAD encounters any errors, the P_CLI system call returns to the calling process with the error code set by the P_LOAD system call.

A successful load operation establishes the command file in memory with its Base Page partially initialized. The P_CLI system call then continues parsing the command tail to set up the Base Page values from 050H to 0FFH.

P_CLI initializes an unused Process Descriptor from the internal PD table, a UDA (expanded UDA if 8087 processing is required) and a 96-byte stack area. The UDA and stack are dynamically allocated from memory. P_CLI then calls the P_CREATE system call. If P_CLI encounters an error in any of these steps, it releases all memory segments allocated for the new command, as well as the Process Descriptor, and then returns with the appropriate error code set.

Once the P_CREATE system call returns successfully, the P_CLI system call assigns the calling process's default virtual console to the new process and then returns.

The calling process should set its priority to less than the TMP (198) if it wants to attach to the virtual console after the created process releases it. Once the calling process has successfully reattached, it should set its priority back to 200.

Refer to Table 6-5 for a list of error codes returned in CX.

P_CREATE	
Create A Process	
Entry Parameters:	
Register CL:	090H (144)
DX:	PD Address - Offset
DS:	PD Address - Segment PD filled in
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The P_CREATE system call allows a process to create a subprocess within its own memory area. The child process shares all memory owned by the calling process at the time of the P_CREATE call. If the Process Descriptor (PD) is outside of the operating system area, the system copies it into a PD from the internal PD Table. The system call returns an error code if there are no more unused PDs in the table.

The User Data Area (UDA) can be anywhere in memory but is required to be on a paragraph boundary. The only time the system copies the PD is if it is not within 64k of the System Data Segment.

Process Descriptors, as well as Queue Descriptors and Queue Buffers, are required to be within the System Data Segment because they are linked together on various system lists or are used by more than one process. Because of this, they cannot be in the Transient Process Area (TPA), where they cannot be protected.

More than one process can be created by a single P_CREATE call if the LINK field of the PD is nonzero. In this case, it is assumed to point to another PD within the same Data Segment. After it creates the first process, the system call checks the Process Descriptor's LINK field. Using this linked list of PDs, a single P_CREATE call can create multiple processes.

Note: The P_CREATE system call does not check the validity of the PD addresses passed by the calling process. An invalid PD address can cause Concurrent CP/M to crash if no hardware memory protection is available on the system.

Refer to Table 6-5 for a list of error codes returned in CX.

00	LINK	THREAD	STAT	PRIOR	FLAG
08	NAME				
10	UDA	DISK	USER	RESERVED	MEM
18	RESERVED				PARENT
20	CNS	RESERVED	LIST	RESERVED	SFLAG
28	RESERVED				

Figure 6-12. PD - Process Descriptor

Table 6-17. PD Field Definitions

<i>Field</i>	<i>Definition</i>
LINK	Link field for insertion on current system list. If this field's initial value is nonzero, it is assumed to point to another PD. This field is used to create more than one process with a single Create Process call.
THREAD	Link field for insertion on Thread List. Initialized to be zero (0).
STAT	Current Process activity. Initialized to be zero (0). Activity codes are listed below:
00 RUN	The process is ready to run. The STAT field is always in this state when a process is examining its own Process Descriptor. The PD is on the Ready List. The currently running process is always at the head of Ready List.
01 POLL	The process is polling a device. The PD is on the Poll List.
02 DELAY	The process is delaying for a specified number of system ticks. The PD is on the Delay List.
06 Read Queue	The process is waiting to read a message from a system queue that is empty. The PD is on the Read Queue List whose root is in the Queue Descriptor of the system queue involved.
07 Write Queue	The process is waiting to write a message to a system queue whose buffer is full. The PD is on the Write Queue List, whose root is in the Queue Descriptor of the system queue involved.

Table 6-17. (continued)

<i>Field</i>	<i>Definition</i>
08 FLAGWAIT	The process is waiting for a system flag to be set. The PD is in the flag table entry of the flag it is waiting for.
09 CIOWAIT	The process is waiting to attach to a character I/O device (console or list) while another process owns it. The PD is on CQUEUE list whose root is in the Character Control Block of the device in question.
PRIOR	<p>Current priority. Process scheduling is done based on this field. Typical user programs run at a priority of 200. 0 is the best priority, and 255 is the worst priority. The following is a list of priority assignments used by most Concurrent CP/M systems. User processes priorities should be from 200-254.</p> <ul style="list-style-type: none"> 1 Initialization Process 2-31 Interrupt Handlers 32-63 System Processes 64-190 Undefined 191-197 Undefined 198 Terminal Message Process 199 Undefined 200 Default Priority For Transients 201-254 User Processes 255 Idle Process
FLAG	<p>Bit field of flags determining run-time characteristics of a process. Initialize as needed. All undocumented flags are used internally or are reserved for system use.</p> <p>001H SYS System Process. Has privileged access to various features of Concurrent CP/M. This process can only be terminated if the termination code is 0FFH. This process can access restricted system queues. This flag is turned off if the calling process is not a system process.</p>

Table 6-17. (continued)

<i>Field</i>	<i>Definition</i>
002H KEEP	This process cannot be terminated. This flag is turned off if the calling process is not a system process.
004H KERNEL	This process resides within the operating system. This flag is turned off if the PD is not within the operating system.
010H TABLE	This PD is copied into the PD from the PD table. When this process terminates, the PD is recycled into the PD table.
8000H 8087	This process is an 8087-running process.
NAME	Process Name. Eight bytes, all eight bits of each byte are used for matching process names.
UDA	Segment address of this process's User Data Area. Initialized to be the number of paragraphs from the beginning of the calling process's Data Segment. The User Data Area contains process information that is not needed between processes. It also contains the System Stack of each process. Refer to the UDA description below.
DISK	Current default disk
USER	Current default user number
MEM	Root of linked list of Memory Segment Descriptors that are owned by this process. Initialized to zero, except for reentrant or shared code RSPs.
SFLAG	Second Flag. If bit 0 of SFLAG (01H) is set, the system suspends this process whenever it is switched out to the background and runs it only when it is switched in to the foreground.
PARENT	Process that created this process. The P_CREATE system call sets this value at process creation. The parent field is set to zero if the parent terminates before the child.

Table 6-17. (continued)

<i>Field</i>	<i>Definition</i>
CNS	Current default console's number. Initialized to be the default console number.
LIST	Current default list device's number. Initialized to be the default list device number.
RESERVED	Reserved for internal use. These fields must be initialized to zero (0).




Figure 6-13. UDA - User Data Area

The length of the UDA is 256 bytes (352 bytes if 8087 processing is required), and it must begin on a paragraph boundary.

Table 6-18. UDA Field Definition

<i>Field</i>	<i>Definition</i>												
DMA OFFS	The initial DMA OFFSET for the new process. The segment address of the DMA is assumed to be the same as the initial Data Segment (refer to DS below)												
AX,BX,CX,DX, DI,SI,BP	The initial register values for the new process. These are typically set to zero.												
SP	<p>The initial stack pointer for the new process. The stack pointer is relative to the initial Stack Segment (refer to SS below). The initial stack of the new process must be initialized with the offset of the first instruction to be executed by the new process. The word that the stack pointer points to is the initial instruction pointer. Two words must follow the initial IP, which is filled in with the initial Code Segment (refer to CS below) and the initial flags. The initial flags are set to 0200H, which means that interrupts are on, and all other flags are off. Concurrent CP/M starts a new process by executing an Interrupt Return instruction with the initial stack.</p> <p>Note: This stack area is distinct from the User System Stack at the end of the UDA.</p> <div style="text-align: center; margin-top: 20px;"> <p>Low Memory</p> <p>·</p> <p>stack area</p> <p>·</p> <table style="margin-left: auto; margin-right: auto;"> <tr> <td style="padding-right: 20px;">SS</td> <td style="padding-right: 20px;">SP</td> <td style="padding-right: 20px;">IP</td> <td></td> </tr> <tr> <td></td> <td></td> <td style="text-align: center;">0</td> <td style="padding-left: 20px;">(CS)</td> </tr> <tr> <td></td> <td></td> <td style="text-align: center;">0</td> <td style="padding-left: 20px;">(Flags)</td> </tr> </table> <p style="margin-top: 20px;">Stack Initialization Area</p> </div>	SS	SP	IP				0	(CS)			0	(Flags)
SS	SP	IP											
		0	(CS)										
		0	(Flags)										

Table 6-18. (continued)

<i>Field</i>	<i>Definition</i>
INT 0, INT 1, INT 3, INT 4	The initial interrupt vectors for the first five interrupt types can be set by filling in these fields. The first word of each field is the Instruction Pointer (IP), and the second word is the Code Segment (CS) for a list of the interrupt routine that services these interrupts. Those fields that are zero are initialized to be the same as the calling processes interrupt vectors. These fields are typically initialized to be 0.
CS,DS, ES,SS	The initial segment addresses for the new process are taken from these fields. Those fields that are zero are initialized to be the same as the calling process's Data Segment.
INT 224, INT 225	Interrupts 224 and 225 are used to communicate with Concurrent CP/M by typical programs. These interrupt vectors are initialized to be the same as the calling process if these values are zero. The ability to change these values allows a run-time system to intercept Concurrent CP/M calls that its children make. The suggested protocol is to keep INT 225 pointing to the Concurrent CP/M entry point and changing INT 224 to point to an internal routine. When a child process does an INT 224, the internal routine can filter calls to Concurrent CP/M using INT 225 for the actual Concurrent CP/M call.
RESERVED	All reserved fields are used internally and must be initialized to zero.
USER SYSTEM STACK	This is the stack area used by the process when it is in the operating system. The SP variable in the UDA should not point to this area.
CW*	Control word for 8087 processor. Processes bypassing the P_CLI or P_LOAD system call must set this word to 03FFH.
SW*	Status word for 8087 processor. Processes bypassing the P_CLI or P_LOAD system call must set this word to 0000H.
*Part of optional 8087 Extension. If the 8087 flag is set in the SFLAG field, this 6-paragraph extension must be included for the 8087 environment.	

P_DELAY
Delay For Specified Number Of System Ticks
Entry Parameters: Register CL: 08DH (141) DX: Number of System Ticks

The P_DELAY system call causes the calling process to wait until the specified number of system ticks has occurred. The P_DELAY system call avoids the necessity of programmed delay loops. It allows other processes to use the CPU resource while the calling process waits.

The length of the system tick varies among installations. A typical system tick is 60Hz (16.67 milliseconds). In Europe, it is likely to be 50Hz (20 milliseconds). The exact length of the system tick can be obtained by reading the TICKS/SEC value from the System Data Segment (refer to the S_SYSDAT system call).

There is up to one tick of uncertainty in the exact amount of time delayed. This is due to the P_DELAY system call being called asynchronously from the actual time base. The P_DELAY system call is guaranteed to delay the calling process at least the number of ticks specified. However, when the calling process is rescheduled to run, it might wait quite a bit longer if there are higher priority processes waiting to run. The P_DELAY system call is used primarily by programs that need to wait specific amounts of time for I/O events to occur. Under these conditions, the calling process usually has a very high priority level. If a process with a high priority calls the P_DELAY system call, the actual delay is typically within a system tick of the amount of time wanted.

P_DISPATCH
Call Dispatcher
Entry Parameters: Register CL: 08EH (142)

The P_DISPATCH system call forces a reschedule of processes that are waiting to run. Normally, dispatches occur at every system tick interrupt (usually 60 times a second), and whenever a process releases a system resource. Dispatching also occurs whenever a process needs a system resource that is not currently available. A CPU-bound process runs for no more than one system tick before a dispatch is forced. The dispatch occurs at the next system tick.

The Concurrent CP/M Dispatcher is priority driven, with round-robin scheduling of equivalent-priority processes. When a process calls the P_DISPATCH system call, it is rescheduled, so that processes with higher or equivalent priorities are given the CPU before the calling process obtains it again. The calling process regains control of the CPU resource when it becomes the highest priority process again.

P_LOAD	
Load A CMD Type File Into Memory	
Entry Parameters:	
Register CL:	03BH (59)
DX:	FCB Address - Offset
DS:	FCB Address - Segment
Returned Values:	
Register AX:	Base Page Address 0FFFFH on error
BX:	Same as AX
CX:	Error Code

The P_LOAD system call loads a disk CMD type file into memory. Upon entry, register DX contains the offset, relative to DS, of a successfully opened FCB that specifies the CMD file to load. Upon return, register AX has the value 0FFFFH if the program load failed. Otherwise, AX contains the paragraph address of the Base Page belonging to the loaded program. The paragraph address and length of each group loaded from the CMD file is found in the Base Page. See Sections 3.2 and 3.3.

Note that before calling P_LOAD, the calling process must establish the DMA address of where the CMD file is to be loaded. This is accomplished with F_DMASEG and F_DMAOFF.

Note: Open the CMD file in Read-Only mode and close it once the load is completed.

Refer to Table 6-5 for a list of error codes returned in CX.

P_PDADR
Return The Address Of The Calling Process's Process Descriptor
Entry Parameters: Register CL: 09CH (156)
Returned Values: Register AX: PD Address - Offset BX: Same as AX ES: PD Address - Segment

The P_PDADR system call obtains the address of the calling process's Process Descriptor. For a description of the format of the Process Descriptor, refer to the P_CREATE system call.

P_PRIORITY
Set The Priority Of The Calling Process
Entry Parameters: Register CL: 091H (145) DL: Priority

The P_PRIORITY system call sets the priority of the calling process to the specified value. This system call is useful in situations where a process needs to have a high priority during an initialization phase, but afterwards can run at a lower priority.

The best or highest priority is 00H, while the worst or lowest priority is 0FFH. Transient processes are initialized to run at C8H (200 decimal) by the P_CLI system call.

P_RPL	
Resident Procedure Library	
Entry Parameters:	Register CL: 097H (151) DX: CPB Address-Offset DS: CPB Address - Segment
Returned Values:	Register AX: 01H if RPL not found RPL return parameter BX: same as AX CX: Error Code ES: RPL return segment if addr




Figure 6-14. CPB - Call Parameter Block

Table 6-19. CPB Field Definitions

<i>Field</i>	<i>Definition</i>
NAME	Name of Resident Procedure, eight ASCII characters
PARAM	Parameter to send to the Resident Procedure

P_RPL permits a process to call a system call in an optional Resident Procedure Library (RPL).

P_RPL opens a system queue with the specified name. If the Q_OPEN system call succeeds, P_RPL checks the queue to verify that it is an RPL-type queue. If either the Q_OPEN fails, or if it is not an RPL-type queue, P_RPL returns to the calling process with an error code.

P_RPL reads a message from the queue that contains the address of the specified system call. It then places the PARAM field of the CPB in register DX, and places the calling process's Data Segment address in register DS. P_RPL performs a Far Call instruction to the address it obtains from the queue message. Upon return from the RPL, the system call copies the BX register to the AX register and then returns to the calling process.

Note: The P_RPL system call does not write the address of the Resident Procedure back to the queue. The Resident Procedure itself must do this. If the Resident Procedure is to be reentrant, it must write the message into the queue upon entry. If it is to be serially reusable, the procedure must write the message just before returning.

Refer to Table 6-5 for a list of error codes returned in CX.

P_TERM
Terminate Calling Process
<p>Entry Parameters:</p> <p style="padding-left: 40px;">Register CL: 08FH (143)</p> <p style="padding-left: 40px;">DL: Term Code</p> <p>Returned Values:</p> <p style="padding-left: 40px;">Register AX: 0FFFFH on failure</p> <p style="padding-left: 40px;">BX: Same as AX</p> <p style="padding-left: 40px;">CX: Error Code</p>

The P_TERM system call terminates the calling process. If the termination code is not 0FFH, the system call can only terminate a user process. If the termination code is 0FFH, the system call can terminate the calling process even though the process's SYSTEM flag is on. P_TERM cannot terminate a process with the KEEP flag on. If the termination is successful, the system call releases the mutual exclusion queues owned by the process. It also releases all memory segments owned by the process, and returns the Process Descriptor to the PD table.

A process can own one or more of the following resources: memory segments, consoles, printers, mutual exclusion messages, and system Lock List entries that record open files and locked records. When a process terminates and releases its resources, these resources become available to other processes on the system. For example, if a terminating process releases a system console, the console is usually given back to the console's TMP. This occurs when the TMP is the highest priority process waiting for the console.

If the system call returns to the calling process, the P_TERM call has failed for one of two reasons. Either the process has the KEEP flag on, or it has the SYSTEM flag on, and the termination code is not 0FFH.

P_TERMCPM	
Entry Parameters:	
Register CL:	00H (0)
Returned Values:	
Register AX:	0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The P_TERMCPM system call terminates the calling process, releasing all system resources owned by the process.

P_TERMCPM is implemented internally by calling P_TERM with the termination code set to 00H.

Under CP/M-86, the P_TERMCPM system call has a further argument that allows a process not to release its memory. This argument places a piece of code into memory that becomes an interface for later programs. Concurrent CP/M does not include this option. Memory segments are not recovered by the system until all processes that own the memory segment have released it.

Refer to Table 6-5 for a list of returned error codes.

6.2.8 Queue System Calls

Queue system calls under Concurrent CP/M use the Queue Parameter Block data structure to pass parameters to and from the operating system. Listing 6-3 shows the structure of the Queue Parameter Block and the equates for its fields.




Figure 6-15. QPB - Queue Parameter Block

Table 6-20. QPB Field Definitions

<i>Field</i>	<i>Description</i>
QUEUEID	Queue number field; filled in by a Q_OPEN operation
* 0000H	Reserved for internal use; must be initialized to zero
BUFFER	Offset address of Queue Message Buffer
NAME	Name of Queue for Q_OPEN operation

```

;*****
;*
;*      QPB - Queue Parameter Block Definition
;*
;*      -
;* 00      0000H      queueid      0000H      buffer
;*      -
;* 08              name
;*      -
;*
;*      queueid - Queue ID, address of QD
;*      buffer  - address to read/write into/from
;*      name    - name of queue (for open only)
;*
;*****

qpb_0          equ      word ptr 0
qpb_queueid   equ      word ptr qpb_0 + word
qpb_buffer    equ      word ptr qpb_queueid + 4
qpb_name      equ      byte ptr qpb_buffer + word

qpb_len       equ      qpb_name + qnamsiz
qnamsiz       equ      8

```

Listing 6-3. Queue Parameter Block Definition

Q_CREAD	
Conditionally Read A Message From A System Queue	
Entry Parameters:	
Register CL:	08AH (138)
DX:	QPB Address - Offset
DS:	QPB Address - Segment
	QPB_queueid filled in by previous Q_OPEN
	QPB_buffer set to message buffer offset
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code message in buffer

The Q_CREAD system call is analogous to the Q_READ system call, but it returns an error code if there are not enough messages to read, instead of waiting for another process to write to the queue.

Refer to Table 6-5 for a list of error codes returned in CX.

Q_CWRITE	
Conditionally Write A Message To A System Queue	
Entry Parameters:	
Register CL:	08CH (140)
DX:	QPB Address - Offset
DS:	QPB Address - Segment
	QPB_queueid filled in by previous Q_OPEN
	QPB_buffer set to message buffer offset message in current DMA buffer
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code

The Q_CWRITE system call is analogous to the Q_WRITE system call, but it returns an error code if there is not enough system queue buffer space for the message to be written, instead of waiting for another process to read from the queue.

Refer to Table 6-5 for a list of error codes returned in CX.

Q_DELETE	
Delete A System Queue	
<p>Entry Parameters:</p> <p>Register CL: 088H (136)</p> <p>DX: QPB Address - Offset</p> <p>DS: QPB Address - Segment</p> <p>QPB_queueid filled in by a previous Q_OPEN call</p>	
<p>Returned Values:</p> <p>Register AX: 0 on success 0FFFFH on failure</p> <p>BX: Same as AX</p> <p>CX: Error Code</p>	

The Q_DELETE system call removes a system queue from the system. The system returns error codes if the queue cannot be deleted or if the queue has not been opened prior to the Q_DELETE call.

Refer to Table 6-5 for a list of error codes returned in CX.

Q_MAKE	
Make A System Queue	
<p>Entry Parameters:</p> <p style="margin-left: 20px;">Register CL: 086H (134)</p> <p style="margin-left: 40px;">DX: QD Address - Offset</p> <p style="margin-left: 40px;">DS: QD Address - Segment</p> <p style="margin-left: 60px;">QD filled in</p>	
<p>Returned Values:</p> <p style="margin-left: 20px;">Register AX: 0 on success</p> <p style="margin-left: 40px;">0FFFFH on failure</p> <p style="margin-left: 40px;">BX: Same as AX</p> <p style="margin-left: 40px;">CX: Error Code</p>	




Figure 6-16. QD - Queue Descriptor

Table 6-21. Queue Descriptor Field Definitions

<i>Field</i>	<i>Definition</i>
FLAGS	<p>Queue Flags. The bits are defined as follows</p> <p>0001H - Mutual exclusion queue 0002H - Cannot be deleted 0004H - Restricted to system processes 0008H - RSP message queue 0010H - Used internally 0020H - RPL address queue 0040H - Used internally 0080H - Used internally</p> <p>Remaining flags reserved for future use</p>
NAME	8-byte queue name. All 8 bits of each character are matched on a Q_OPEN call.
MSGLEN	Number of bytes in each logical message
NMSGs	Maximum number of logical messages to be supported. If the number of messages written to the queue equals this maximum, no more messages are allowed until a message is read.
BUFFER	Address of the queue buffer. This buffer must be (NMSGs * MSGLEN) bytes long. The address is an offset relative to the DS register. This field is unused if the QD resides outside of the System Data Segment. Typically this field is 0 if the queue is being created by a transient program. RSPs that create queues must initialize this field to point to a buffer. The Data Segment of an RSP's queue is considered part of the System Data Segment unless it is beyond 64k of the beginning of the System Data Segment.
* 0000H	For internal use. Must be initialized to zero.

Every system queue under Concurrent CP/M is associated with a Queue Descriptor that resides within the Concurrent CP/M System Data Segment. In the Q_MAKE system call, the calling process passes the address of a Queue Descriptor. If this Queue Descriptor is within the Concurrent CP/M System Data Segment, the system uses it directly for the System Queue. If the Queue Descriptor is outside of the System Data Segment, the system obtains a Queue Descriptor from an internal Queue Descriptor table. If there are no unused Queue Descriptors in the internal table, the system call returns an error code.

Refer to Table 6-5 for a list of error codes returned in CX.

The buffer for a system queue must also reside within the System Data Area. For non-00H length buffers, resident buffers are used directly. The system obtains a buffer from the Queue Buffer Area if the buffer does not reside within the System Data Segment. The size of the buffer is calculated from the NMSGs and MSGLEN fields. The system call returns an error code if there is not enough unused buffer area left to accommodate this new buffer.

All system queues must have unique names. The system call returns an error code if a system queue already exists by the given name.

Under Concurrent CP/M, all system queues must be explicitly opened (refer to the Q_OPEN system call) before being used to read or write messages or to delete the queue.

Q_OPEN	
Open A System Queue	
Entry Parameters:	
Register CL:	087H (135)
DX:	QPB Address - Offset
DS:	QPB Address - Segment
	QPB_name filled in
Returned Values:	
Register AX:	0 on success 0FFFFH on failure
BX:	Same as AX
CX:	Error Code
	QPB_queueid filled in

All system queues under Concurrent CP/M must be explicitly opened before a read, write, or delete operation can be done. The Q_OPEN system call examines each existing system queue and attempts to match the name in the QPB with the name of a system queue. All eight bytes of the name must match for a successful open. All bits of each byte are examined. If the open operation is successful, the Q_OPEN system call modifies the Queue ID Field of the QPB. Once the the queue is opened, subsequent reads, writes, or a delete are allowed.

Refer to Table 6-5 for a list of error codes returned in CX.

Q_READ	
Read A Message From A System Queue	
Entry Parameters:	
Register CL:	089H (137)
DX:	QPB Address - Offset
DS:	QPB Address - Segment
	QPB_queueid filled in by previous Q_OPEN
	QPB_buffer set to message buffer offset
Returned Values:	
Register AX:	0 on success
	0FFFFH on failure
BX:	Same as AX
CX:	Error Code message in buffer

The Q_READ system call reads a message from a system queue that was previously opened by the calling process. The system call returns an error code if the queue was not previously opened or if the system queue has been deleted since the Q_OPEN call. If there are not enough messages to read from the queue, the calling process waits until another process writes into the queue before returning.

Refer to Table 6-5 for a list of error codes returned in CX.

Q_WRITE	
Write A Message To A System Queue	
<p>Entry Parameters:</p> <p>Register CL: 08BH (139)</p> <p>DX: QPB Address - Offset</p> <p>DS: QPB Address - Segment</p> <p>QPB_queued filled in by previous Q_OPEN</p> <p>QPB_buffer set to message buffer offset</p>	
<p>Returned Values:</p> <p>Register AX: 0 on success 0FFFFH on failure</p> <p>BX: Same as AX</p> <p>CX: Error Code</p>	

The Q_WRITE system call writes a message to a system queue that was previously opened by the calling process. The system call returns an error code if the queue was not previously opened or if the system queue has been deleted since the Q_OPEN call. If there is not enough buffer space in the queue, the calling process waits until another process reads from the queue before writing to the queue and returning.

Refer to Table 6-5 for a list of error codes returned in CX.

6.2.9 System Information System Calls

S_BDOSVER	
Return BDOS Version Number	
Entry Parameters:	
Register CL:	0CH (12)
Returned Values:	
Register AL:	31 (BDOS Version 3.1)
AH:	14 (Concurrent CP/M)
BX:	Same as AX

The S_BDOSVER system call returns the BDOS file system version number, allowing version-independent programming.

AL High Nibble = BDOS Version Number	
AL Low Nibble = BDOS Revision Level	
AH High Nibble = CPU Type	
0 = 8080	
1 = 8086	
AH Low Nibble = OS Type	
0 = CP/M	2 = CP/M w/networking
1 = MP/M	3 = MP/M w/networking
4 = Concurrent CP/M	6 = Concurrent CP/M
5,7 to E = Reserved	w/networking

Figure 6-17. BDOS Version Number Format

S_BIOS	
Call BIOS Character Routine	
<p>Entry Parameters:</p> <p style="margin-left: 40px;">Register CL: 032H (50)</p> <p style="margin-left: 80px;">DX: BIOS Desc. Addr. - Offset</p> <p style="margin-left: 80px;">DS: BIOS Desc. Addr. - Segment</p> <p>Returned Values:</p> <p style="margin-left: 40px;">Register AX: BIOS Return</p> <p style="margin-left: 80px;">BX: Same as AX</p>	




Figure 6-18. BIOS Descriptor Format

The S_BIOS system call is provided under Concurrent CP/M for compatibility with programs generated under CP/M-86 that use this system call (Function 50). Under Concurrent CP/M, only routines that interface with character devices are supported. The arguments to character routines such as CONIN and LIST must be converted to those appropriate for the Concurrent CP/M XIOS. Refer to the *Concurrent CP/M System Guide* for further information about the XIOS.

Note: Calls to the XIOS Console Status, Input, and Output system calls do not go to the XIOS if the referenced device is a virtual console.

S_OSVER	
Return The Version Of Current Concurrent CP/M System	
Entry Parameters:	
Register CL:	0A3H (163)
Returned Values:	
Register AX:	Version Number (01431H)
BX:	Same as AX
CX:	Error Code

The S_OSVER system call provides information that allows version-independent programming. The system call returns a two-byte value, with AH set to 014H for Concurrent CP/M, and AL set to the Concurrent CP/M version level. The AH register contains a value set to the type of operating system. A value of 01431H indicates Concurrent CP/M 3.1.

Refer to Table 6-5 for a list of error codes returned in CX.

AL High Nibble = Concurrent CP/M-86 Version Number	
AL Low Nibble = Concurrent CP/M Revision Level	
AH High Nibble = CPU Type	
0 = 8080	
1 = 8086	
AH Low Nibble = OS Type	
0 = CP/M	2 = CP/M w/networking
1 = MP/M	3 = MP/M w/networking
4 = Concurrent CP/M	6 = Concurrent CP/M
5,7 to E = Reserved	w/networking

Figure 6-19. Operating System Version Number Format

S_SERIAL
Return Current System's Serial Number
<p>Entry Parameters:</p> <p style="margin-left: 40px;">Register CL: 06BH (107)</p> <p style="margin-left: 40px;">DX: SERIAL Address - Offset</p> <p style="margin-left: 40px;">DS: SERIAL Address - Segment</p> <p>Returned Values:</p> <p style="margin-left: 40px;">SERIAL filled in</p>




Figure 6-20. SERIAL Number Format

S_SERIAL returns the Concurrent CP/M serial number to the addressed, six-byte SERIAL field as a six-byte ASCII numeral.

S_SYSDAT	
Return Address Of The System Data Segment	
Entry Parameters:	
Register CL:	09AH (154)
Returned Values:	
Register AX:	Sysdat Address - Offset
Register BX:	Same as AX
Register ES:	Sysdat Address - Segment

The S_SYSDAT system call returns the address of the System Data Segment of the calling process. The System Data Segment contains all Process Descriptors, Queue Descriptors, the roots of system lists, and other internal data that Concurrent CP/M uses.

Figure 6-21, illustrates the SYSDAT Table and its fields.

00H	SUP ENTRY				RESERVED			
08H	RESERVED							
10H	RESERVED							
18H	RESERVED							
20H	RESERVED							
28H	XIOS ENTRY				XIOS INIT			
30H	RESERVED							
38H	DISPATCHER				PDISP			
40H	CCPMSEG		RSPSEG		ENDSEG		RESER- VED	NVCNS
48H	NLCB	NCCB	N- FLAGS	SYS- DISK	MMP		RESER- VED	DAY FILE
50H	TEMP DISK	TICKS /SEC	LUL		CCB		FLAGS	
58H	MDUL		MFL		PUL		QUL	
60H	QMAU							
68H	RLR		DLR		DRL		PLR	
70H	RESERVED		THRDRT		QLR		MAL	
78H	VERSION		VERNUM		CCPMVERNUM		TOD_DAY	
80H	TOD _HR	TOD _MIN	TOD _SEC	NCON DEV	NLST DEV	NCIO DEV	LCB	
88H	OPEN_FILE		LOCK_ MAX	OPEN_ MAX	OWNER_8087		RESERVED	
90H	RESERVED							
98H	RESERVED							XPCNS
A0H	OFF_8087		SEG_8087		RESERVED			

Figure 6-21. SYSDAT Table

Table 6-22. SYSDAT Table Data Fields

<i>Field</i>	<i>Explanation</i>
SUP ENTRY	Double-word address of the Supervisor entry point for inter-module communication. All internal system calls go through this entry point.
XIOS ENTRY	Double-word address of the Extended I/O System entry point for intermodule communication. All XIOS function calls go through this entry point.
XIOS INIT	Double-word address of the Extended I/O System Initialization entry point. System hardware initialization takes place by a call through this entry point.
DISPATCHER	Double-word address of the Dispatcher entry point that handles interrupt returns. Executing a Far Jump to this address is equivalent to executing an Interrupt Return instruction. The Dispatcher routine causes a dispatch to occur and then executes an Interrupt Return. All registers are preserved and one level of stack is used. This location should be used as an exit point by all XIOS interrupt handlers that use the DEV_SETFLAG system call.
PDISP	Double-word address of the Dispatcher entry point that causes a dispatch to occur with all registers preserved. Once the dispatch is done, a RETF instruction is executed. Executing a JMPF PDISP is equivalent to executing a RETF instruction. This location should be used as an exit point whenever the XIOS releases a resource that might be wanted by a waiting process.
CCPMSEG	Starting paragraph of the operating system area. This is also the Code Segment of the Supervisor Module.
RSPSEG	Paragraph Address of the first RSP in a linked list of RSP Data Segments. The first word of the data segment points to the next RSP in the list. Once the system has been initialized, this field is zero.

Table 6-22. (continued)

<i>Field</i>	<i>Explanation</i>
ENDSEG	First paragraph beyond the end of the operating system area, including any buffers consisting of uninitialized RAM allocated to the operating system by GENCCPM. These include the Directory Hashing, Disk Data and XIOS ALLOC buffers. These buffer areas, however, are not part of the CCPM.SYS file.
NVCNS	Number of virtual consoles, copied from the XIOS Header by GENCCPM.
NLCB	Number of List Control Blocks, copied from the XIOS Header by GENCCPM.
NCCB	Number of Character Control Blocks, copied from the XIOS Header by GENCCPM.
NFLAGS	Number of system flags as specified during GENCCPM.
SYSDISK	Default system disk. The CLI looks on this disk if it cannot open the command file on the user's current default disk. Set during GENCCPM.
MMP	Maximum memory allowed per process. Set during GENCCPM.
DAY FILE	Day File option. If this field is OFFH, the operating system displays file logging information on system consoles at each command. Set during GENCCPM.
TEMP DISK	Default temporary disk. Programs that create temporary files should use this disk. Set during GENCCPM.
TICKS/SEC	The number of system ticks per second.
LUL	Link list root of unused Lock List items.
CCB	Address of the Character Control Block Table, copied from the XIOS Header by GENCCPM.

Table 6-22. (continued)

<i>Field</i>	<i>Explanation</i>
FLAGS	Address of the Flag Table.
MDUL	Link list root of unused Memory Descriptors.
MFL	Link list root of free memory partitions.
PUL	Link list root of unused Process Descriptors.
QUL	Link list root of unused Queue Descriptors.
QMAU	Queue Buffer Memory Allocation Unit.
RLR	Ready List Root. Linked list of PDs that are ready to run.
DLR	Delay List Root. Link list of PDs that are delaying for a specified number of system ticks.
DRL	Dispatcher Ready List. Temporary holding place for PDs that have just been made ready to run.
PLR	Poll List Root. Linked list of PDs that are polling on devices.
THRDRT	Thread List Root. Linked list of all current PDs on the system. The list is threaded through the THREAD field of the PD instead of the LINK field.
QLR	Queue List Root. Linked list of all System QDs.
MAL	Link list of active memory allocation units. A MAU is created from one or more memory partitions.
VERSION	Address, relative to CCPMSEG, of version string.
VERNUM	Concurrent CP/M version number (system call 12, S_BDOSVER).
CCPMVERNUM	Concurrent CP/M version number (system call 163, S_OSVER).

Table 6-22. (continued)

<i>Field</i>	<i>Explanation</i>
TOD_DAY	Time-of-Day. Number of days since 12/31/77.
TOD_HR	Time-of-Day. Hour of the day.
TOD_MIN	Time-of-Day. Minute of the hour.
TOD_SEC	Time-of-Day. Second of the minute.
NCONDEV	Number of XIOS consoles, copied from the XIOS Header by GENCCPM.
NLSTDEV	Number of XIOS list devices, copied from the XIOS Header by GENCCPM.
NCIODEV	Total number of character devices (NCONDEV + NLSTDEV).
LCB	Offset of the List Control Block Table, copied from the XIOS Header by GENCCPM.
OPEN_FILE	Open File Drive Vector. Designates drives that have open files on them. Each bit of the word value represents a disk drive; the least significant bit represents Drive A, and so on through the most significant bit, Drive P. Bits which are set indicate drives containing open files.
LOCK_MAX	Maximum number of locked records per process. Set during GENCCPM.
OPEN_MAX	Maximum number of open disk files per process. Set during GENCCPM.
OWNER_8087	Specifies 8087 information. If set to 0FFFFH, the system assumes there is no 8087 in the system. If set to 0, there is an 8087 but no one owns it. If set to any other value, the system assumes that this value is the PD offset of the 8087 current process.

Table 6-22. (continued)

<i>Field</i>	<i>Explanation</i>
XPCNS	Specifies the number of physical consoles.
OFF_8087	Offset of the hardware-dependent 8087 interrupt vector. If you supply your own 8087 exception handler routine, store the offset of your exception handler routine at this offset address.
SEG_8087	Segment address of the hardware-dependent 8087 interrupt vector. If you supply your own 8087 exception handler routine, store the segment address of your exception handler routine at this segment address.

T_GET	
Get System Time And Date	
<p>Entry Parameters:</p> <p style="padding-left: 20px;">Register CL: 069H (105)</p> <p style="padding-left: 40px;">DX: TOD Address - Offset</p> <p style="padding-left: 40px;">DS: TOD Address - Segment</p>	
<p>Returned Values:</p> <p style="padding-left: 20px;">Register AL: Seconds</p> <p style="padding-left: 40px;">TOD filled in</p> <p style="padding-left: 40px;">(Days, Hours and Minutes only)</p>	



Figure 6-22. TOD - Time-of-Day Structure

Table 6-23. Time-of-Day Field Definitions

<i>Field</i>	<i>Definition</i>
DAY	The number of days since 12/31/77. The day is stored as a 16-bit integer.
HOUR	The current hour of the current day. The hour is represented as a 24 hour clock in 2 binary coded decimal (BCD) digits.
MIN	The current minute of the current hour. The minute is stored as 2 BCD digits.
SEC	The current second of the current minute. The second is stored as 2 BCD digits.

The T_GET system call obtains the system internal time and date. The calling process passes the address of a four-byte data structure that receives the time and date values. This system call is equivalent to the T_SECONDS system call, except that it does not return the SECONDS field of the internal time.

T_SECONDS	
Get Current System Time And Day	
Entry Parameters:	
Register CL:	09BH (155)
DX:	TOD Address - Offset
DS:	TOD Address - Segment
Returned Values:	
	TOD filled in (Days, Hours, Minutes, and Seconds)

The T_SECONDS system call returns the current encoded time and date (including seconds) in the TOD structure passed by the calling process.

T_SET
Set System Time And Date
Entry Parameters: Register CL: 068H (104) DX: TOD Address - Offset DS: TOD Address - Segment

The T_SET system call sets the system internal time and date. The calling process passes the address of a 4-byte structure containing the time and date specification.

The date is represented as a 16-bit integer with day 1 corresponding to January 1, 1978. The time is represented as two bytes hours and minutes stored as two BCD digits.

Under Concurrent CP/M, this system call also sets the second field of the system time and date to 00H.

End of Section 6

Appendix A

System Call Summary by Function Number

This appendix lists the Concurrent CP/M system calls by function number including the parameters a process must pass when calling the function, and the values the function returns to the process.

Table A-1. System Call Summary by Function Number

<i>Dec</i>	<i>Hex</i>	<i>Mnemonic</i>	<i>Input Parameters</i>	<i>Returned Values</i>
0	0	P_TERMCPM	none	AX = Rtn Code
1	1	C_READ	none	AL = char
2	2	C_WRITE	DL = char	none
5	5	L_WRITE	DL = char	none
6	6	C_RAWIO	see def	see def
9	9	C_WRITESTR	DX = .Buffer	none
10	A	C_READSTR	DX = .Buffer	see def
11	B	C_STAT	none	AL = 1 if ready = 0 if not ready
12	C	S_BDOSVER	none	AX = Version#
13	D	DRV_ALLRESET	none	see def
14	E	DRV_SET	AL = Drive #	see def
15	F	F_OPEN	DX = .FCB	AL = Dir Code
16	10	F_CLOSE	DX = .FCB	AL = Dir Code
17	11	F_SFIRST	DX = .FCB	AL = Dir Code
18	12	F_SNEXT	none	AL = Dir Code
19	13	F_DELETE	DX = .FCB	AL = Dir Code
20	14	F_READ	DX = .FCB	AL = Err Code
21	15	F_WRITE	DX = .FCB	AL = Err Code
22	16	F_MAKE	DX = .FCB	AL = Dir Code
23	17	F_RENAME	DX = .FCB	AL = Dir Code
24	18	DRV_LOGINVEC	none	AX = Login Vect.
25	19	DRV_GET	none	AL = Cur Drive
26	1A	F_DMAOFF	DX = .DMA	none
27	1B	DRV_ALLOCVEC	none	ES:AX = Alloc Addr

Table A-1. (continued)

<i>Dec</i>	<i>Hex</i>	<i>Mnemonic</i>	<i>Input Parameters</i>	<i>Returned Values</i>
28	1C	DRV_SETRO	none	see def
29	1D	DRV_ROVEC	none	AX = R/O Vect.
30	1E	F_ATTRIB	DX = .FCB	see def
31	1F	DRV_DPB	none	ES:AX = DPB Addr
32	20	F_USERNUM	DL = 0FFH (get) = User # (set)	AL = User # none
33	21	F_READRAND	DX = .FCB	AL = Err Code
34	22	F_WRITERAND	DX = .FCB	AL = Err Code
35	23	F_SIZE	DX = .FCB	R0, R1, R2 AL = Dir Code
36	24	F_RANDREC	DX = .FCB	R0, R1, R2
37	25	DRV_RESET	DX = drive Vect	AL = Err Code
38	26	DRV_ACCESS	DX = drive Vect	none
39	27	DRV_FREE	DX = drive Vect	none
40	28	F_WRITEZF	DX = .FCB	AL = Err Code
42	2A	F_LOCK	DX = .FCB	AL = Err Code
43	2B	F_UNLOCK	DX = .FCB	AL = Err Code
44	2C	F_MULTISEC	DL = # of Records	AL = Rtn Code
45	2D	F_ERRMODE	DL = Error Mode	none
46	2E	DRV_SPACE	DL = Drive #	see def
47	2F	P_CHAIN	see def	none
48	30	DRV_FLUSH	none	see def
50	32	S_BIOS	DX = .BD	AX = BIOS Rtn
51	33	F_DMASEG	DX = .DMA Seg	none
52	34	F_DMAGET	none	ES:AX = DMA Addr
53	35	MC_MAX	DX = .MCB	see def
54	36	MC_ABSMAX	DX = .MCB	see def
55	37	MC_ALLOC	DX = .MCB	see def
56	38	MC_ABSALLOC	DX = .MCB	see def
57	39	MC_FREE	DX = .MCB	see def
58	3A	MC_ALLFREE	none	none
59	3B	P_LOAD	DX = .FCB	AX = BP Addr
99	63	F_TRUNCATE	DX = .FCB	see def
100	64	DRV_SETLABEL	DX = .FCB	AL = Dir Code
101	65	DRV_GETLABEL	DX = Drive #	AL = Label Data Byte
102	66	F_TIMEDATE	DX = .XFCB	AL = Dir Code
103	67	F_WRITEXFCB	DX = .XFCB	AL = Dir Code
104	68	T_SET	DX = .TOD	none

Table A-1. (continued)

<i>Dec</i>	<i>Hex</i>	<i>Mnemonic</i>	<i>Input Parameters</i>	<i>Returned Values</i>
105	69	T_GET	DX = .TOD	AL = seconds
106	6A	F_PASSWD	DX = .Password	none
107	6B	S_SERIAL	DX = .serial#	serial #
109	6D	C_MODE	DX = Con Mode = 0FFFFH	none AX = Con Mode
110	6E	C_DELIMIT	DL = Out Delim = 0FFFFH	none AL = Out Delim
111	6F	C_WRITEBLK	DX = .CHCB	none
112	70	L_WRITEBLK	DX = .CHCB	none
128	80	M_ALLOC	DX = .MPB	AX = Rtn Code
129	81	M_ALLOC	Same as above	Same as above
130	82	M_FREE	DX = .MPB	none
131	83	DEV_POLL	DL = Device	none
132	84	DEV_WAITFLAG	DL = Flag	AX = Rtn Code
133	85	DEV_SETFLAG	DL = Flag	AX = Rtn Code
134	86	Q_MAKE	DX = .QD	none
135	87	Q_OPEN	DX = .QPB	AX = Rtn Code
136	88	Q_DELETE	DX = .QPB	AX = Rtn Code
137	89	Q_READ	DX = .QPB	none
138	8A	Q_CREAD	DX = .QPB	AX = Rtn Code
139	8B	Q_WRITE	DX = .QPB	AX = Rtn Code
140	8C	Q_CWRITE	DX = .QPB	AX = Rtn Code
141	8D	P_DELAY	DX = #ticks	none
142	8E	P_DISPATCH	none	none
143	8F	P_TERM	DL = Term. Code	AX = Rtn Code
144	90	P_CREATE	DX = .PD	none
145	91	P_PRIORITY	DL = Priority	none
146	92	C_ATTACH	none	none
147	93	C_DETACH	none	none
148	94	C_SET	DL = Console	none
149	95	C_ASSIGN	DX = .ACB	AX = Rtn Code
150	96	P_CLI	DX = .CLBUF	none
151	97	P_RPL	DX = .CPB	AX = result
152	98	F_PARSE	DX = .PFCB	see def
153	99	C_GET	none	AL = con #
154	9A	S_SYSDAT	none	ES:AX = Sys Data Addr
155	9B	T_SECONDS	DX = .TOD	TOD filled in
156	9C	P_PDADR	none	ES:AX = PD Addr

Table A-1. (continued)

<i>Dec</i>	<i>Hex</i>	<i>Mnemonic</i>	<i>Input Parameters</i>	<i>Returned Values</i>
157	9D	P_ABORT	DX = .ABP	AX = Rtn Code
158	9E	L_ATTACH	none	none
159	9F	L_DETACH	none	none
160	A0	L_SET	DL = List #	none
161	A1	L_CATTACH	none	AX = Rtn Code
162	A2	C_CATTACH	none	AX = Rtn Code
163	A3	S_OSVER	none	AX = Version #
164	A4	L_GET	none	AL = List #

Conventions used in Appendix A:

.	= Address of
#	= Number
ACB	= Assign Control Block
Addr	= Address
APB	= Abort Parameter Block
BD	= Bios Descriptor
BP	= Base Page
Char	= ASCII Character
CHCB	= Character Control Block
CLBUF	= Command Line Buffer
Con	= Console
CPB	= Call Parameter Block
Cur	= Current
Delim	= Delimiter
Dir	= Directory
DMA	= Direct Memory Address
Err	= Error
FCB	= File Control Block
MCB	= Memory Control Block
MPB	= Memory Parameter Block
Num	= Number
Out	= Output
PD	= Process Descriptor
PFCB	= Parse Filename Control Block
QD	= Queue Descriptor
QPB	= Queue Parameter Block
Rec	= Record
Rtn	= Return
Sys	= System
Term.	= Termination
TOD	= Time of Day
Vect	= Vector

End of Appendix A

Appendix B

ASCII and Hexadecimal Conversions

This appendix contains tables of the ASCII symbols, including their binary, decimal, and hexadecimal conversions.

Table B-1. ASCII Symbols

<i>Symbol</i>	<i>Meaning</i>	<i>Symbol</i>	<i>Meaning</i>
ACK	acknowledge	FS	file separator
BEL	bell	GS	group separator
BS	backspace	HT	horizontal tabulation
CAN	cancel	LF	line feed
CR	carriage return	NAK	negative acknowledge
DC	device control	NUL	null
DEL	delete	RS	record separator
DLE	data link escape	SI	shift in
EM	end of medium	SO	shift out
ENQ	enquiry	SOH	start of heading
EOT	end of transmission	SP	space
ESC	escape	STX	start of text
ETB	end of transmission	SUB	substitute
ETX	end of text	SYN	synchronous idle
FF	form feed	US	unit separator
		VT	vertical tabulation

Table B-2. ASCII Conversion Table

<i>Binary</i>	<i>Decimal</i>	<i>Hexadecimal</i>	<i>ASCII</i>
0000000	000	00	NUL
0000001	001	01	SOH (CTRL-A)
0000010	002	02	STX (CTRL-B)
0000011	003	03	ETX (CTRL-C)
0000100	004	04	EOT (CTRL-D)
0000101	005	05	ENQ (CTRL-E)

Table B-2. (continued)

<i>Binary</i>	<i>Decimal</i>	<i>Hexadecimal</i>	<i>ASCII</i>
0000110	006	06	ACK (CTRL-F)
0000111	007	07	BEL (CTRL-G)
0001000	008	08	BS (CTRL-H)
0001001	009	09	HT (CTRL-I)
0001010	010	0A	LF (CTRL-J)
0001011	011	0B	VT (CTRL-K)
0001100	012	0C	FF (CTRL-L)
0001101	013	0D	CR (CTRL-M)
0001110	014	0E	SO (CTRL-N)
0001111	015	0F	SI (CTRL-O)
0010000	016	10	DLE (CTRL-P)
0010001	017	11	DC1 (CTRL-Q)
0010010	018	12	DC2 (CTRL-R)
0010011	019	13	DC3 (CTRL-S)
0010100	020	14	DC4 (CTRL-T)
0010101	021	15	NAK (CTRL-U)
0010110	022	16	SYN (CTRL-V)
0010111	023	17	ETB (CTRL-W)
0011000	024	18	CAN (CTRL-X)
0011001	025	19	EM (CTRL-Y)
0011010	026	1A	SUB (CTRL-Z)
0011011	027	1B	ESC (CTRL-[)
0011100	028	1C	FS (CTRL-)
0011101	029	1D	GS (CTRL-])
0011110	030	1E	RS (CTRL-^)
0011111	031	1F	US (CTRL-_)
0100000	032	20	(SPACE)
0100001	033	21	!
0100010	034	22	"
0100011	035	23	#
0100100	036	24	\$
0100101	037	25	%
0100110	038	26	&
0100111	039	27	'
0101000	040	28	(
0101001	041	29)
0101010	042	2A	*
0101011	043	2B	+

Table B-2. (continued)

<i>Binary</i>	<i>Decimal</i>	<i>Hexadecimal</i>	<i>ASCII</i>
0101100	044	2C	,
0101101	045	2D	-
0101110	046	2E	.
0101111	047	2F	/
0110000	048	30	0
0110001	049	31	1
0110010	050	32	2
0110011	051	33	3
0110100	052	34	4
0110101	053	35	5
0110110	054	36	6
0110111	055	37	7
0111000	056	38	8
0111001	057	39	9
0111010	058	3A	:
0111011	059	3B	;
0111100	060	3C	<
0111101	061	3D	=
0111110	062	3E	>
0111111	063	3F	?
1000000	064	40	@
1000001	065	41	A
1000010	066	42	B
1000011	067	43	C
1000100	068	44	D
1000101	069	45	E
1000110	070	46	F
1000111	071	47	G
1001000	072	48	H
1001001	073	49	I
1001010	074	4A	J
1001011	075	4B	K
1001100	076	4C	L
1001101	077	4D	M
1001110	078	4E	N
1001111	079	4F	O
1010000	080	50	P

Table B-2. (continued)

<i>Binary</i>	<i>Decimal</i>	<i>Hexadecimal</i>	<i>ASCII</i>
1010001	081	51	Q
1010010	082	52	R
1010011	083	53	S
1010100	084	54	T
1010101	085	55	U
1010110	086	56	V
1010111	087	57	W
1011000	088	58	X
1011001	089	59	Y
1011010	090	5A	Z
1011011	091	5B	[
1011100	092	5C	\
1011101	093	5D]
1011110	094	5E	^
1011111	095	5F	<
1100000	096	60	'
1100001	097	61	a
1100010	098	62	b
1100011	099	63	c
1100100	100	64	d
1100101	101	65	e
1100110	102	66	f
1100111	103	67	g
1101000	104	68	h
1101001	105	69	i
1101010	106	6A	j
1101011	107	6B	k
1101100	108	6C	l
1101101	109	6D	m
1101110	110	6E	n
1101111	111	6F	o
1110000	112	70	p
1110001	113	71	q
1110010	114	72	r
1110011	115	73	s
1110100	116	74	t
1110101	117	75	u

Table B-2. (continued)

<i>Binary</i>	<i>Decimal</i>	<i>Hexadecimal</i>	<i>ASCII</i>
1110110	118	76	v
1110111	119	77	w
1111000	120	78	x
1111001	121	79	y
1111010	122	7A	z
1111011	123	7B	{
1111100	124	7C	
1111101	125	7D	}
1111110	126	7E	~
1111111	127	7F	DEL

End of Appendix B

Appendix C

Error Codes

Table C-1. Concurrent CP/M Error Codes

<i>Code #</i>	<i>Definition</i>
0	NO ERROR
1	FUNCTION NOT IMPLEMENTED
2	ILLEGAL FUNCTION NUMBER
3	CAN'T FIND MEMORY
4	ILLEGAL SYSTEM FLAG NUMBER
5	FLAG OVERRUN
6	FLAG UNDERRUN
7	NO UNUSED QUEUE DESCRIPTORS LEFT IN QD TABLE
8	NO UNUSED QUEUE BUFFER AREA LEFT
9	CAN'T FIND QUEUE
10	QUEUE IN USE
12	NO UNUSED PROCESS DESCRIPTORS LEFT IN PROCESS DESCRIPTOR TABLE
13	QUEUE ACCESS DENIED
14	EMPTY QUEUE
15	FULL QUEUE
16	CLI QUEUE MISSING
17	NO 8087 IN SYSTEM
18	NO UNUSED MEMORY DESCRIPTORS LEFT IN MEMORY DESCRIPTOR TABLE
19	ILLEGAL CONSOLE NUMBER
20	CAN'T FIND PROCESS DESCRIPTOR BY NAME
21	CONSOLE DOES NOT MATCH
22	NO CLI PROCESS
23	ILLEGAL DISK NUMBER
24	ILLEGAL FILE NAME
25	ILLEGAL FILE TYPE
26	CHARACTER NOT READY
27	ILLEGAL MEMORY DESCRIPTOR
28	BAD LOAD
29	BAD READ
30	BAD OPEN

Table C-1. (continued)

<i>Code #</i>	<i>Definition</i>
31	NULL COMMAND
32	NOT OWNER
33	NO CODE SEGMENT IN LOAD FILE
34	ACTIVE PROCESS DESCRIPTOR
35	CAN'T TERMINATE
36	CAN'T ATTACH
37	ILLEGAL LIST DEVICE NUMBER
38	ILLEGAL PASSWORD
40	EXTERNAL TERMINATION OCCURRED
41	FIXUP ERROR UPON LOAD
42	FLAG SET IGNORED

End of Appendix C

Appendix D

ECHO.A86 Listing

Listing D-1. ECHO.A86

```
;
;      ECHO - Resident System Process
;      Print Command tail to console
;
;
;      DEFINITIONS
;
ccpmint          equ      224      ;ccpm entry interrupt
c_writestr       equ      9        ;print string
c_detach         equ      147      ;detach console
c_set            equ      148      ;set default console
q_make          equ      134      ;create queue
q_open           equ      135      ;open queue
q_read           equ      137      ;read queue
q_write         equ      139      ;write queue
p_priority       equ      145      ;set priority

pdlen            equ      48        ;length of Process
;      Descriptor

p_cns            equ      byte ptr 020h ;default cns
p_disk           equ      byte ptr 012h ;default disk
p_user           equ      byte ptr 013h ;default user
p_list           equ      byte ptr 024h ;default list
ps_run           equ      0         ;PD run status
pf_keep         equ      2         ;PD nokill flag

rsp_top          equ      0         ;rsp offset
rsp_pd           equ      010h      ;PD offset
rsp_uda          equ      040h      ;UDA offset
rsp_bottom       equ      140h      ;end rsp header

qf_rsp           equ      08h       ;queue RSP flag
```

Listing D-1. (continued)

```

;
; CODE SEGMENT
;
CSEG
org 0

ccpm:  int ccpmint
      ret

main:  ;create ECHO queue
      mov cl,q_make ! mov dx,offset qd
      call ccpm
      ;open ECHO queue
      mov cl,q_open ! mov dx,offset qpb
      call ccpm
      ;set priority to normal
      mov cl,p_priority ! mov dx,200
      call ccpm

      ;ES points to SYSDAT
      mov es,sdatseg

loop:  ;forever
      ;read cmdtail from queue
      mov cl,q_read ! mov dx,offset qpb
      call ccpm

      ;set default values from PD
      mov bx,pdadr
;      mov dl,es:p_disk[bx] ;p_disk=0-15
;      inc dl ! mov disk,dl ;make disk=1-16
;      mov dl,es:p_user[bx]
;      mov user,dl
;      mov dl,es:p_list[bx]
;      mov list,dl
;      mov dl,es:p_cns[bx]
      mov console,dl

```

Listing D-1. (continued)

```

                ;set default console
;   mov dl,console
   mov cl,C_SET ! call ccpm

                ;scan cmdtail and look for '$' or 0.
                ;when found, replace w/ cr,lf,'$'

   lea bx,cmdtail ! mov al,'$' ! mov ah,0
   mov dx,bx ! add dx,131
nextchar:
   cmp bx,dx ! ja endcmd
   cmp [bx],al ! je endcmd
   cmp [bx],ah ! je endcmd
   inc bx ! jmps nextchar
endcmd:
   mov byte ptr [bx],13
   mov byte ptr 1[bx],10
   mov byte ptr 2[bx],'$'

                ;write command tail

   lea dx,cmdtail ! mov cl,C_WRITESTR
   call ccpm

                ;detach console
   mov dl,console
   mov cl,c_detach ! call ccpm
                ;done, get next command
   jmps loop

;
;   DATA SEGMENT
;

```

Listing D-1. (continued)

```

DSEG
org    rsp_top

sdatseg    dw    0,0,0
           dw    0,0,0
           dw    0,0

           org    rsp_pd

pd         dw    0,0           ; link,thread
           db    ps_run       ; status
           db    190          ; priority
           dw    pf_keep      ; flags
           db    'ECHO'       ; name
           dw    offset uda/10h ; uda seg
           db    0,0          ; disk,user
           db    0,0          ; load dsk,usr
           dw    0            ; mem
           dw    0,0          ; dvract,wait
           db    0,0
           dw    0
           db    0            ; console
           db    0,0,0
           db    0            ; list
           db    0,0,0
           dw    0,0,0,0

```

Listing D-1. (continued)

```

org     rsp_uda

uda     dw     0,offset dma,0,0      ;0
        dw     0,0,0,0
        dw     0,0,0,0              ;10h
        dw     0,0,0,0
        dw     0,0,0,0              ;20h
        dw     0,0,0,0
        dw     0,0,offset stack_tos,0 ;30h
        dw     0,0,0,0
        dw     0,0,0,0              ;40h
        dw     0,0,0,0
        dw     0,0,0,0              ;50h
        dw     0,0,0,0
        dw     0,0,0,0              ;60h

org     rsp_bottom

qbuf    rb     131                  ;Queue buffer

qd      dw     0                    ;link
        db     0,0                  ;net,org
        dw     qf_rsp               ;flags
        db     'ECHO'               ;name
        dw     131                  ;msglen
        dw     1                    ;nmsgs
        dw     0,0                  ;dq,nq
        dw     0,0                  ;msgcnt,msgout
        dw     offset qbuf          ;buffer addr.
    
```

Listing D-1. (continued)

```

dma                rb        128

stack              dw        0ccccch,0ccccch,0ccccch
                  dw        0ccccch,0ccccch,0ccccch
                  dw        0ccccch,0ccccch,0ccccch
                  dw        0ccccch,0ccccch,0ccccch
                  dw        0ccccch,0ccccch,0ccccch
stack_tos          dw        offset main      ; start offset
                  dw        0                ; start seg
                  dw        0                ; init flags

pdadr              rw        1                ; QPB Buffer
cmdtail           rb        129              ; starts here
                  db        13,10,'$'

qpb               db        0,0              ;must be zero
                  dw        0                ;queue ID
                  dw        1                ;nmsgs
                  dw        offset pdadr     ;buffer addr.
                  db        'ECHO '         ;name to open

console           db        0
;disk             db        0
;user            db        0
;list            db        0

end

```

End of Appendix D

Appendix E

8087 Exception Handling

This appendix includes an example of an 8087 interrupt handling routine to demonstrate the requirements for using the 8087 processor. Refer to Intel's *iAPX 86,88 User's Manual* for a description of 8087 exception handling in the section on "8087 Numeric Data Processor".

In order to guarantee the data integrity for each 8087 process in the multitasking environment, any user-defined exception handler must adhere to a minimum sequence of steps within the exception handler:

1. Save the 8086 environment of the 8086-running process.
2. Save the environment of the 8087-running process. The OWNER_8087 field in SYSDAT will contain the offset of the 8087-running process (see description of SYSDAT in Section 6 with the S_SYSDAT system call).
3. Clear the 8087 interrupt request bit in the status word.
4. Disable the 8087 interrupts.
5. Clear the PIC interrupt (this instruction is hardware-dependent).
6. At this point, you might want to modify the 8087 environment image saved in step 2 above.
7. Before enabling the 8086 interrupts, restore the 8087 environment with its status word's interrupt request bit cleared. If the environment is not restored before 8086 interrupts are enabled, and an interrupt occurs (like a tick), a different 8087 process can gain control of the 8087 and swap in its 8087 context. On a second interrupt, or on an IRET instruction, the 8086-running process that happened to be executing the exception handler code is brought back into 8086 context and writes over the new 8087 context.

The user program, which uses its own exception handler, must replace the system's interrupt vector with its own. Once this is done, the system swaps this vector into memory every time the program comes back into 8087 context. The address of the interrupt vector is in the SYSDAT table at offset A0H (the description of the SYSDAT Table is included in the description of the S_SYSDAT system call in Section 6).

The default exception handler aborts those 8087 programs that have enabled 8087 interrupts and that generate a severe error (such as stack underrun, divide by zero, and so forth). Any other errors are ignored by the default exception handler.

Listing E-1. 8087 Exception Handling

```
;=====
ndpint:      ; 8087 interrupt routine
;=====
;           This exception handler is non-specific and
;           is meant as an example
;           default. It is assumed that if the 8087
;           programmer has enabled 8087
;           interrupts and has specified exception flags
;           in the control word, then
;           the programmer has also included an
;           exception handler to take
;           specific actions within the program
;           before continuing in the 8087.
;           This handler will ignore non-severe
;           errors (overflow,etc) and will
;           terminate processes with severe errors
;           (divide by zero,stack violation).
```

Listing E-1. (continued)

```

push ds                ; SAVE CURRENT DATA SEGMENT
mov ds,sysdat          ; GET XIOS DATA SEGMENT
mov ndp_ssreg,ss      ; DO STACK SWITCH FOR 8086 ENVIRONMENT
mov ndp_spreg,sp      ; SAVE
mov ss,sysdat
mov sp,offset ndp_tos ; SAVE THE 8086 REGISTERS
push ax! push bx
push cx! push dx
push di! push si
push bp! push es
mov es,sysdat          ; NOW SAVE THE 8087 ENVIRONMENT
FNSTENV env_8087      ; SAVE 8087 PROCESS INFO
FWAIT
FNCLEX                 ; CLEAR ITS INT REQUEST BIT
xor ax,ax
FNDISI                ; DISABLE ITS INTERRUPTS
mov al,020h           ; SEND 2 INTERRUPT ACKNOWLEDGES - 1 FOR
out 060h,al           ; ONE FOR MASTER PIC, ONE FOR SLAVE
mov al,020h
out 058h,al           ; IN 8087 WILL CHECK THE 8087 ERROR
call in_8087          ; CONDITION. IF ERROR IS SEVERE, IT
                        ; WILL ABORT, ELSE IT WILL RETURN WITH
                        ; NO CHANGES.
mov bx,offset env_8087 ; CLEAR ITS STATUS WORD FOR ENV RESTORE
mov byte ptr 2[bx],0
pop es! pop bp        ; RESTORE THE 8086 ENVIRONMENT
pop si! pop di
pop dx! pop cx
pop bx! pop ax

mov ss,ndp_ssreg      ; SWITCH BACK TO PREVIOUS STACK
mov sp,ndp_spreg
FLDENV env_8087       ; RESTORE 8087 ENV WITH GOOD STATUS
FWAIT
pop ds                ; RESTORE PREVIOUS DATA SEGMENT
iret

```

Listing E-1. (continued)

```

;
;
in_8087:
;-----
;   entry: DS = SYSDAT
;   Only user-specified error conditions generate
;   interrupts from the 8087.

mov bx,owner_8087      ; GET THE PROCESS DESCRIPTOR
test bx,bx             ; CHECK IF OWNER HAS ALREADY
jz end_87              ; TERMINATED
mov si,offset env_8087 ; IF IT'S A SEVERE ERROR, TERMINATE
mov ax,statusw[si]    ;
                       ; IF NOT SEVERE, RETURN & CONTINUE
test ax,03ah          ; 3A = UNDER/OVERFLOW, PRECISION,
jnz end_87            ;      AND DENORMALIZED OPERAND
or p_flag[bx],080h    ; NOT 3A = ZERO DIVIDE OR INVALID
                       ; OPERATION (STACK ERROR)
end_87:
ret

```

End of Appendix E

Glossary

Base Page: Memory region between 0000H and 0100H relative to the beginning of the Data Segment used to hold system parameters. Base Page serves primarily as an interface region between user programs. Note that in the 8080 Model, the code and data are intermixed in the code segment.

BCD: Acronym for Binary Coded Decimal. Representation of decimal numbers using binary digits. See Table B-2 for representations of ASCII codes.

BDOS: Basic Disk Operating System (BDOS). The BDOS manages the Concurrent CP/M file structure and executes most of the Concurrent CP/M system calls.

block: Basic unit of disk space allocation under Concurrent CP/M. Each disk drive has a fixed block size (BLS) defined in its disk Parameter Block in the XIOS. The block size can be 1K, 2K, 4K, 8K, or 16K of consecutive bytes. Blocks are numbered relative to zero on a disk. Blocks are not shared between files.

Boolean: Variable that can have only two values; usually interpreted as true/false or on/off.

Checksum Vector (CSV): Contiguous data area in the XIOS with one byte for each directory sector to be checked, that is, CKS bytes. A Checksum Vector is initialized and maintained for each logged-in drive. Each directory access by the system results in a checksum calculation that is compared with that in the Checksum Vector. If there is a discrepancy, the drive is set to Read-Only status. This prevents the user from inadvertently switching disks without logging in the new disk with a CTRL-C. If not logged in, the new disk is treated the same as the old one, and you can destroy data on it if you write to it.

CIO: Character I/O (CIO) Module. The CIO module handles all character I/O to and from consoles and list devices.

CLI: Command Line Interpreter. The P_CLI system call interprets the command requested in a command line and performs the system calls needed to open a process, load the command file, and execute the code.

CMD: Filetype for Concurrent CP/M command files. These are machine language object modules ready to be loaded and executed. Any file with this type can be executed by simply typing the filename after the drive prompt. For example, the program PIP.CMD can be executed by simply typing PIP.

command: Set of instructions that are executed when the command name is typed after the system prompt. These instructions can be built in the Concurrent CP/M system or can reside on disk as a file of type CMD. Concurrent CP/M commands consist of three parts: the command name, the command tail, and a carriage return.

console: Primary I/O device used by Concurrent CP/M. The console usually consists of a CRT screen for displaying output and a keyboard for input.

control character: Nonprinting ASCII character produced on the console by holding down the CTRL (CONTROL) key while striking the character key. CTRL-H means hold down CTRL and press H. Control characters are sometimes indicated using the up-arrow symbol (^), so CTRL-H can be represented as ^H. Certain control characters are treated as special commands by Concurrent CP/M.

Default Buffer: 128-byte buffer maintained at 0080H in the Base Page. When the CLI loads a CMD file, it initializes this buffer to the command tail, that is, any characters typed after the CMD file name. The first byte at 0080H contains the length of the command tail while the command tail itself begins at 0081H. A binary zero terminates the command tail value. The I command under DDT[™] initializes this buffer in the same way as the CLI.

Default FCB: One of two FCBs maintained at 005CH and 006CH in the Base Page. The P_CLI system call initializes the first default FCB from the first delimited field in the command tail and initializes the second default FCB from the next field in the command tail.

delimiters: ASCII characters used to separate constituent parts of a file specification. The P_CLI system call recognizes certain delimiter characters as : . = ; <> _ ' blank, and carriage return. Several Concurrent CP/M commands also treat ; [] () , and \$ as delimiter characters. It is advisable to avoid the use of delimiter characters and lowercase characters in filenames.

directory: Portion of a disk containing entries for each file on the disk and locations of the blocks allocated to the files. Each file directory entry is in the form of a 32-byte FCB, although one file can have several entries, depending on its size. The maximum number of directory entries supported is specified in the drive's Disk Parameter Block.

directory entry: 32-byte entry associated with each disk file. A file can have more than one directory entry associated with it. There are four directory entries per directory sector. Directory entries can also be referred to as directory FCBs.

disk, diskette: Magnetic media used for mass storage of data in the computer system. The term disk can refer to a diskette, a removable cartridge disk, or a fixed hard disk.

Disk Parameter Block (DPB): Table residing in the XIOS that defines the characteristics of a drive in the disk subsystem used with Concurrent CP/M. The address of the DPB is in the Disk Parameter Header at DPbase + 0AH. Drives with the same characteristics can use the same DPB. However, each logical drive must have its own Disk Parameter Header and DPB. The address of the drive's Disk Parameter Header must be returned in registers HL when the BDOS calls the SELDSK entry point in the XIOS. DRV_DPB returns the DPB address.

Disk Parameter Header (DPH): 16-byte area in the XIOS containing information about the disk drive and a scratchpad area for certain BDOS operations. See the *Concurrent CP/M System Guide* for further details.

extent (EX): 16K consecutive bytes in a file. Extents are numbered from 0 to 31. One extent can contain 1, 2, 4, 8, or 16 blocks. EX is the extent number field of an FCB and is a one-byte field at FCB + 12, where FCB labels the first byte in the FCB. Depending on the Block Size (BLS) and the maximum data Block Number (DSM), a directory entry contains 1, 2, 4, 8, or 16 extents. The EX field is usually set to 0 by the user, but contains the current extent number during file I/O. The term "Extent Folding" describes directory entries containing more than one extent. In CP/M version 1.4, each FCB contained only one extent.

FCB: See File Control Block.

file: Collection of data containing from zero to 242,144 records. Each record contains 128 bytes and can contain either binary or ASCII data. Files consist of one or more 16K extents, with 128 records per extent.

File Control Block (FCB): Thirty-six consecutive bytes maintained and updated by system calls for file I/O. The FCB fields are described in Section 2.4.

hex file format: Absolute output of ASM86 for the Intel 8086. A HEX file contains a sequence of absolute records, which give a load address and byte values to be stored starting at the load address (refer to Section 4.3).

I/O: Acronym for Input/Output operations or routines handling the input and output of data in the computer system.

logical drive: Logically distinct region of a physical drive. A physical drive can be divided into one or more logical drives, and designated with specific drive references (such as a: or f:). Thus, at the user interface, it appears that there are several disks in the system.

MEM: Memory Module. The Memory Module handles all memory management calls by methods transparent to your applications program.

parse: Separate a command line into its syntactic parts.

queue: Data structure used by the file system to keep track system information, such as processes ready to run, locked files, and resources currently in use by processes. Processes also use queues to communicate with one another. The BDOS system calls create and maintain queues.

Read-Only: Condition in which a logical disk drive can be read but not written to. A drive can be set to Read-Only status by using the SET utility. This protects the user from switching disks without executing a disk reset. Files can also be set to Read-Only status with the SET utility or the F_ATTRIB system call. Read-Only is often abbreviated as R/O.

record: Smallest unit of data in a disk file that can be read or written. A record consists of 128 consecutive bytes whose byte displacement in a file is the product of the Record Number times 128. A 128-byte record in a file occupies one 128-byte sector on the diskette. If the blocking and deblocking algorithm is used, several records can occupy each disk sector.

reentrant code: Code that can be used by one process while another is already executing it. Reentrant code must not be self-modifying; it must be pure code that does not contain data. The data for reentrant code can be kept in a separate data area or placed on the stack.

RSP: Reserved System Process. An RSP is a Concurrent CP/M utility included within Concurrent CP/M during the execution of GENCCPM.

RTM: Real Time Monitor. The RTM is the nucleus of Concurrent CP/M, managing queues and flags, polling devices, and dispatching and suspending processes. Application programs gain access to RTM functions through system calls.

sector: Unit of data read from and written to the disk by the XIOS. The sector size is dependent on the disk drive hardware and is usually a power of two, such as 256, 512, 1024, or 2048 bytes. These disk sectors are referred to as Host Sectors.

source file: ASCII text file usually created with a text editor that is an input file to a program, such as a compiler, assembler, or a text formatter.

stack: Reserved area of memory where the processor saves the return address when it receives a Call instruction. When the processor encounters a Return instruction, it restores the current address on the stack to the Instruction Pointer. Data such as the contents of the registers can also be saved on the stack on a first-in-last-out basis. The Push instruction places data on the stack and the Pop instruction removes it. 8086 stacks are 16 bits wide; instructions operating on the stack add and remove stack items one word at a time. An item is pushed onto the stack by decrementing the stack pointer (SP) by 2 and writing the item at the SP address. In other words, the stack grows downward in memory.

SUP: The Supervisor (SUP) manages communications between processes and the operating system kernel, and between other operating system modules. All system calls are intercepted by the SUP.

track: Concentric ring on the disk; the standard IBM single density disks have 77 tracks. Each track consists of a fixed number of numbered sectors. Tracks are numbered from 0 to one less than the number of tracks on the disk. Data on the disk media is accessed by combinations of track and sector numbers.

TMP: Terminal Message Processes. The TMPs are Resident System Processes that intercept command lines from the virtual consoles, check for errors, and pass on executable requests to the CLI. The TMP prints the prompt and some system error messages on your console. Each virtual console has an independent TMP heading defining the console's environment, including the default disk, user number, printer, and console.

transient command file: File of type .CMD stored on disk. Such files must be loaded into the system each time they are executed, and therefore execute more slowly than Resident System Processes (RSPs), which are an integral part of the operating system and execute rapidly. Transient commands are created with the GENCMD utility; RSPs are included in the operating system during execution of GENCCPM.

user: Logically distinct subdivision of the directory. Each directory can be divided into 16 user numbers.

wildcard: A ? or * character. The BDOS directory search calls matches ? with any single character and * with multiple characters. Refer to the F_SFIRST and F_SNEXT system calls for further details.

XIOS: Extended I/O System. In Concurrent CP/M, the BDOS is the invariant file-handling system, which operates independent of the hardware implementation. The XIOS is the customizable I/O interface configured for your hardware system by the system manufacturer. The XIOS is similar to the BIOS in CP/M and CP/M-86, but it has been extended to implement virtual consoles and associated features.

End of Glossary

Index

8080 and Small RSP Models, 5-2
8080 keyword, 4-6
8080 Memory Model, 1-12, 3-5, 3-7,
4-1, 4-3, 5-2, 5-6
exception handling, 3-3
8087 Flag
PD, 5-8
processor, 3-2
support, 1-2, 3-2
96-byte initial stack, 3-1
file reference, 2-7

A

absolute address, 4-7
ACB—Assign Control Block
(Figure 6-1), 6-21
access stamp, 2-24, 6-84
address
Flag Table, 6-86
maximum, 4-7
PD, 6-157
queue buffer, 6-169
System Data Segment, 6-178
version string, 6-182
Ahhhh parameter 4-7
AL0, 6-50
AL1, 6-50
Allocation Block Mask, 6-49
Allocation Block Shift Factor, 6-49
allocation vector, 2-39, 6-46
ambiguous reference, 2-6, 6-16
APB—Abort Parameter Block
(Figure 6-10), 6-139
Archive, 6-65
attribute, 2-15

ASM-86 utility, 2-9
asterisk, 2-6
attribute bits, 2-11, 2-14
attribute
compatibility, 2-31
file, 2-14
interface, 2-14
interface F5', 2-30
interface F6', 2-30
AX
UDA field, 6-152
A_Base, 3-4, 5-11

B

B value, 4-7
background, 1-10
backslash, 2-6
backspace, 6-32
BACKSPACE, 6-34
base extent, 6-11, 6-116
Base Page Initialization, 3-5
Base Page, 4-3, 6-141, 6-144
Compact Model, 4-5
initial Data Segment, 3-1
Small Model, 4-4
BASE
MCB, 6-129
Basic Disk Operating System, 1-4,
1-9, 2-1
BDOS, 1-4
BDOS Error Codes, 2-47
BDOS Error mode, 6-45, 6-75
BDOS file system, 2-1
BDOS Multisector Count, 6-113
BDOS physical errors, 2-44

BDOS revision level, 6-174
 BDOS Version Number Format
 (Figure 6-17), 6-174
 BDOS
 Concurrent CP/M, 1-9
 single-tasking CP/M-86, 1-9
 Bhhhh parameter, 4-7
 BIOS, 1-11
 BIOS Descriptor Format
 (Figure 6-18), 6-175
 bit map, 6-56
 BLM, 6-49
 blocking/deblocking, 2-38, 6-52
 BP
 UDA field, 6-152
 BSH, 6-49
 BUFFER field, 5-10
 size, 5-10, 6-73
 BUFFER
 QD field, 6-169
 QPB field, 6-163
 buffers
 disk data, 6-181
 XIOS ALLOC, 6-181
 burst mode, 2-34
 BX
 UDA field, 6-152
 byte count, 2-37, 2-38, 6-65, 6-83

C

C option
 SYSTAT, 1-14
 C(onsole) option, 1-15
 C(onsoles) option
 SYSTAT, 1-14
 C-Seg, 4-11
 Call Parameter Block, 6-159
 carriage return, 2-9, 6-32, 6-33, 6-34,
 6-90

CCB, 1-10
 SYSDAT field, 6-181
 CCPM.SYS file, 5-11, 6-180
 CCPMSEG, 6-182
 SYSDAT field, 6-179
 CCPMVERNUM
 SYSDAT field, 6-182
 Character Control Block, 1-10, 6-39,
 6-128, 6-148, 6-150, 6-181
 character device, 6-175, 6-183
 Character I/O Module, 1-4, 1-10
 CHARACTERS
 C_READSTR, 6-34
 CHCB format, 6-39, 6-128
 checksum, 2-11, 2-17, 2-27, 2-33, 4-12,
 6-68, 6-80, 6-84
 Checksum Vector Size, 6-50
 Checksum Vector Size field
 DPB, 2-40
 checksum verification, 2-27
 disable, 2-33
 child process, 5-10
 CIO, 1-4, 1-10
 CIOWAIT
 Activity code, 6-148
 CKS, 6-50
 CKS field
 Disk Parameter Block, 1-11
 CLBUF, 6-143
 CLI, 1-11, 6-181
 CLI Command Line Buffer
 (Figure 6-11), 6-142
 CLI
 handling RSPs, 5-4
 CLOCK, 1-8
 CLOCK process, 1-2, 1-8
 clock ticks, 1-6
 Close Checksum error, 2-33, 2-45
 CMD, 1-12
 CMD filetype, 6-143

CMD file, 2-9, 4-1, 4-6, 5-5, 6-141, 6-156
 CMD File Header Format (Figure 3-1), 3-3
 CNS
 APB field, 6-140
 C_ASSIGN system call, 6-22
 PD field, 6-150
 Code Group Descriptor, 3-2, 5-2
 Code Segment, 3-2, 6-152, 6-153
 Supervisor, 6-180
 Command Line Buffer, 6-143
 Command Line Interpreter, 1-11, 3-1
 Command RSP, 5-4, 5-5, 5-6
 COMMAND TAIL
 RSP Command Queue Message, 5-5
 COMMAND
 CLI Command Line Buffer, 6-142
 Compact Memory Model, 3-5, 4-5
 Compact Model, 1-12, 4-2, 4-5
 compatibility attribute, 2-15, 2-31
 definition, 2-32, 2-33
 COMPATMODE option
 GENCCPM, 2-32
 compute file size, 2-2
 Concurrent CP/M Compact Memory Model (Figure 4-4), 4-5
 Concurrent CP/M Functional Modules (Figure 1-2), 1-3
 Concurrent CP/M Virtual/Physical Environments (Figure 1-1), 1-1
 Concurrent CP/M Base Page Values (Figure 3-3), 3-6
 concurrent file access, 2-35
 conditional queue write, 6-166
 conditional read
 queue, 1-7
 conditional write
 queue, 1-7
 CONIN, 6-39, 6-175
 CONOUT:, 6-39
 console, 1-11
 Console Buffer Format (Figure 6-2), 6-33
 console I/O, 1-10
 Console I/O System Calls, 6-4, 6-21
 console
 input, 6-131, 6-175
 mode, 6-39
 number, 6-36
 number of XIOS, 6-183
 number of SYSDAT, 6-184
 Output, 6-175
 status, 6-31, 6-175
 system calls, 6-2
 virtual, 6-181
 contiguous memory segment, 6-131
 Continuous display option
 SYSTAT, 1-14, 1-15
 control characters, 2-6
 Control Word
 UDA 8087 extension, 6-153
 copy number
 RSP, 5-3
 CP/M Compatible Memory
 Allocation System Calls, 6-9
 CP/M-86 compatibility, 6-175
 CP/M-86 memory allocation scheme, 6-128
 CPB, 6-160
 CPB—Call Parameter Block (Figure 6-14), 6-159
 CPU type, 6-174, 6-176
 CR byte, 6-113
 CR field, 3-7, 6-79, 6-83, 6-84, 6-93, 6-96
 CR field of FCB, 6-66
 CR field
 FCB, 2-12, 2-38
 CS, 6-153

- CS field
 - FCB, 2-11, 2-38
- CS register
 - Small Model, 4-4
- CS
 - UDA field, 6-153
- CSEG directive
 - ASM-86, 4-4
- CTRL-C, 1-10, 1-15, 5-8, 6-31, 6-32, 6-37
 - disable, 6-29
 - enable, 6-29
- CTRL-E, 6-34
- CTRL-H, 6-34
- CTRL-I, 6-32, 6-34, 6-39
- CTRL-J, 6-34, 6-35
- CTRL-M, 6-34, 6-35
- CTRL-O, 1-10, 1-11, 6-31
 - disable, 6-29
 - enable, 6-29
- CTRL-P, 1-10, 1-11, 6-29, 6-31
- CTRL-Q, 1-10
 - disable, 6-29
 - enable, 6-29
- CTRL-R, 6-35
- CTRL-S, 1-10, 1-11, 6-31
 - disable, 6-29
 - enable, 6-29
- CTRL-U, 6-35
- CTRL-X, 6-35
- CTRL-Z at EOF, 2-9
- current DMA, 6-61, 6-91, 6-96, 6-101
- current DMA address, 6-113
- current DMA buffer, 6-107
- Current Output Delimiter, 6-25
- current processes, 1-13
- current record field, 6-93
 - FCB, 2-12
- current record position, 3-7
- current user number, 2-17, 6-149
- current process activity, 6-147

- CW
 - UDA 8087 extension, 6-153
- CX Error Code Reports, 6-19
- CX error codes, 1-13
- CX
 - UDA field, 6-152
- C_ASSIGN system call, 6-22
- C_ATTACH system call, 6-22, 6-23
- C_CATTACH system call, 6-24
- C_DELIMIT system call, 6-25, 6-40
- C_DETACH system call, 6-22, 6-26
- C_GET system call, 6-27
- C_MODE, 3-1, 6-37
- C_MODE call, 1-11
- C_MODE system call, 6-28
- C_RAWIO, 6-37
- C_RAWIO call, 1-11
- C_RAWIO system call, 6-30
- C_READ system call, 6-32, 6-33, 6-38
- C_READSTR call, 1-10
- C_READSTR system call, 6-33
- C_SET system call, 6-36
- C_STAT, 6-29
- C_STAT system call, 6-37
- C_WRITE, 6-29
- C_WRITE system call, 6-38
- C_WRITEBLK, 6-29
- C_WRITEBLK system call, 6-39
- C_WRITESTR, 6-25, 6-29
- C_WRITESTR system call, 6-40

D

- D0-D15 field
 - FCB, 2-12
- data area, 2-1, 2-8
- data block size, 2-8
- Data Group Descriptor, 5-2, 5-11
- Data Record, 4-9, 4-10
- Data Segment, 5-1

- Data Structures Index, 6-18
- date and time, 1-2
- date and time stamps, 2-3, 2-18, 2-24, 6-61
- DATE utility, 2-25
- Day file option, 6-181
- DAY FILE
 - SYSDAT field, 6-181
- DAY
 - TOD field, 6-185
- days
 - number of, 6-183, 6-185
- DDT-86, 5-11
- Default Close, 6-68
- default console, 6-26, 6-27, 6-150
 - C_ATTACH, 6-23
 - C_CATTACH, 6-24
- default disk, 1-11, 6-47, 6-54, 6-59, 6-149
- default DMA base, 6-74
- default DMA buffer, 3-8, 6-141
- default drive, 2-3, 2-5, 3-7
- default error mode, 1-10, 2-43
- default list device, 6-122, 6-123, 6-124, 6-126, 6-127
- default list device number, 6-125, 6-150
- default mode
 - BDOS Error mode, 6-75
 - Locked mode, 2-26
 - password, 2-3, 2-23, 6-91, 6-107
 - TMP, 2-23
- Delay List, 1-6, 1-9, 6-147, 6-182
- DELAY
 - Activity code, 6-147
- Delete mode, 2-22
- delimiters, 2-6, 6-88
- Device System Calls, 6-2, 6-5, 6-41
- DEV_POLL system call, 1-11, 6-41
- DEV_SETFLAG, 6-42, 6-43, 6-180
- DEV_WAITFLAG, 1-8, 6-42, 6-43

- DI
 - UDA field, 6-152
- DIR attribute, 6-143
- DIR utility, 2-1, 2-15
- Direct Memory Address, 6-73
- direct video mapping, 3-8
- Directory Allocation Vector 0, 6-50
- directory area, 2-1
 - code, 2-46, 2-48, 6-17
 - code definitions, 2-48
 - entry, 6-79
 - label, 2-3, 2-18, 2-19, 2-20, 6-55, 6-60, 6-113
- directory label data byte, 2-19, 2-20, 6-55, 6-60
- Directory Label Format (Figure 2-2), 2-18
- Directory Maximum, 6-50
- Directory Record with SFCB (Figure 2-4), 2-24
- directory space, 2-1
- directory write operations, 2-38
- Disk Data buffers, 6-181
- disk directory area, 2-8
- disk drive organization, 2-8
- Disk Drive System Calls, 6-2, 6-5, 6-44
- Disk File System Calls, 6-7
- Disk Free Space Field Format (Figure 6-5), 6-63
- Disk I/O error, 2-44
- Disk Parameter Block, 1-11, 2-40, 6-48, 6-49
- Disk Reset, 6-51
- Disk Storage Maximum, 6-50
- Disk System Reset (Figure 2-6), 2-41
- DISK
 - DP field, 6-149
- disk
 - temporary, 6-181

- Dispatcher, 1-5
- Dispatcher entry point, 6-180
- Dispatcher Ready List, 6-182
- DISPATCHER
 - SYSDAT field, 6-180
- DL field
 - directory label, 2-19
- DLR
 - SYSDAT field, 6-182
- DMA address, 2-3, 3-1, 6-156
- DMA base, 3-1
- DMA Buffer, 5-6, 5-9, 6-73
- DMA OFFS
 - UDA field, 6-152
- DMA offset, 3-1, 6-72, 6-152
- DMA
 - default address, 6-47
- DPB, 2-40, 6-48
- DPB—Disk Parameter Block
 - (Figure 6-4), 6-48
- DR FCB field, 6-59
- DR field
 - directory label, 2-19
 - FCB, 2-11
 - XFCB, 2-21
- Drive Code
 - FCB, 2-11
- drive code
 - XFCB, 2-21
- drive
 - directory label, 6-101
 - field, 6-89
 - reset, 2-39, 2-41
 - specifier, 2-5
 - status, 2-2, 2-3
- Drive Vector, 6-44
- Drive
 - R/O, or Login Vector Structure (Figure 6-3), 6-44
- DRL
 - SYSDAT field, 6-182
- DRM, 6-50
- DRV_, 2-2
- DRV_ACCESS system call, 2-39, 2-42, 6-45
- DRV_ACCESS call, 2-42
- DRV_ALLOCVEC system call, 6-46
- DRV_ALLRESET, 3-1, 6-73
- DRV_ALLRESET system call, 2-39, 6-47, 6-62
- DRV_DPB system call, 6-48
- DRV_FLUSH system call, 6-52
- DRV_FREE, 2-29
- DRV_FREE system call, 2-39, 2-42, 6-53
- DRV_GET system call, 6-54
- DRV_GETLABEL system call, 2-20 6-55
- DRV_LOGINVEC system call, 6-56
- DRV_RESET, 2-40, 3-1
- DRV_RESET call, 1-11
- DRV_RESET operation, 2-40, 6-62
- DRV_RESET system call, 2-39, 6-57
- DRV_ROVEC system call, 6-47, 6-58
- DRV_SET system call, 6-59
- DRV_SETLABEL system call, 2-19, 6-60
- DRV_SETRO system call, 2-40, 2-42, 2-44, 6-47, 6-58, 6-62
- DRV_SPACE, 6-73
- DRV_SPACE system call, 6-46, 6-64
- DS and ES registers
 - Small Model, 4-4
- DS
 - UDA field, 6-153
- DSEG directive, 4-4
- DSM, 6-50
- DX
 - UDA field, 6-152

E

- E(xit) option, 1-15
 - SYSTAT, 1-14
- ECHO, 5-8
- ECHO RSP, 5-1, 5-3, 5-11
- ENDSEG
 - SYSDAT field, 6-181
- EOF, 6-12
- EOF (CTRL-Z), 2-9
- error codes, 1-13, 2-46, 2-47
- error flag, 2-47, 2-49
- error handling, 2-43
- Error mode, 2-3, 2-43
- ES
 - UDA field, 6-153
- EX field, 6-79
 - FCB, 2-11
- exception handling
 - 8087, 6-184
- exclusive lock, 6-76
- exclusive locks, 2-35
- exit point, 6-180
- EXM, 6-49
- EXT
 - MCB, 6-129
- Extended Address Record, 4-9, 4-10
- extended error codes, 2-49
- Extended Error Module, 1-10
- extended errors, 2-43, 2-45, 2-46
- extended file lock, 2-30, 6-15, 6-107
- Extended I/O System, 1-4
- Extended I/O System entry point, 6-180
- Extended Input/Output System, 1-11
- extent, 6-93
- Extent Mask, 6-49
- extent number
 - FCB, 2-11
- Extra Segments, 5-1

F

- F1' compatibility attribute, 2-32
- F1'-F4', 2-15
- F1'-F4' compatibility attributes, 2-32
- F1'-F4' file attribute, 6-65
- F1'-F8', 2-14
- F1-F8 field
 - FCB, 2-11
- F2' compatibility attribute, 2-33
- F3' attributes, 2-36
- F3' compatibility attribute, 2-33
- F4' compatibility attribute, 2-33
- F5', 2-17
- F5' interface attribute, 2-30, 2-35, 6-65, 6-68
- F5' interface attribute, 6-70
- F5' interface attribute, 6-76, 6-79, 6-107, 6-111
- F5' interface attribute, 2-36
- F5'-F8', 2-16
- F5'-F8' attribute, 6-66
- F6', 2-17
- F6' interface attribute, 2-27, 2-30, 2-36, 2-38, 6-65, 6-68, 6-83
- F7', 2-17
- F8', 2-17
- Far Jump instruction
- Far Return, 3-1, 4-2, 4-3
- FCB, 2-9, 6-17, 6-64

FCB—File Control Block

(Figure 2-1), 2-10

checksum, 2-29

checksum verification, 2-33

drive code, 6-59

extent number, 6-80

format, 2-17

initialization, 2-12

length, 2-10

usage, 2-12

verification, 2-41

FCB

File Name1, 3-7

File Name2, 3-7

FCBADR

PFCB, 6-87

file access, 2-35

concurrent, 2-35

shared, 1-10

File Already Exists error, 2-46

file attributes, 2-14, 6-65

file byte counts, 2-37

File Control Block, 2-9, 6-64

File Currently Open error, 2-45

File field

XFCB, 2-21

file header

CMD, 3-2

File ID, 2-12, 2-26, 2-35, 6-76, 6-80,

6-84, 6-109

File lock, 6-14

extended, 6-65, 6-68

file locking, 1-9

extended, 2-30

file logging information, 6-181

file open modes, 2-26

File Opened in Read/Only Mode

error, 2-45

file

organization, 2-8

security, 2-27

size, 2-8

specification, 2-5

system, 2-1, 2-18, 2-37

system calls, 2-3, 2-4

File-Access System Calls, 6-2, 6-64

filename, 2-1, 6-89

field, 2-1, 2-5

file size,

maximum, 2-8

filetype, 2-1, 6-89

FCB, 2-11

filetype conventions, 2-7

filetype field, 2-5, 2-6, 2-11

XFCB, 2-21

Flag 1

tick flag, 1-9

Flag 2

second flag, 1-8

FLAG field

PD, 6-140

flag IP, 6-42

flag numbers, 6-43

Flag Table

address, 6-182

FLAG

PD field, 6-149

flag

Process Keep, 1-11

SYS, 6-140

flags 0

1, 2, and 3, 6-43

FLAGS field, 5-8, 5-9

flags

initial, 6-152

FLAGS

QD field, 6-169

flags

queue, 6-169

FLAGS

SYSDAT field, 6-182

FLAGWAIT

Activity code, 6-149

flush buffers, 2-39

Function 0, 6-162

Function 1, 6-32

Function 2, 6-38

Function 5, 6-127

Function 6, 6-30

Function 9, 6-40

Function 10, 6-33

Function 11, 6-37

Function 12, 6-174

Function 13, 6-47

Function 14, 6-59

Function 15, 6-83

Function 16, 6-68

Function 17, 6-100

Function 18, 6-104

Function 19, 6-70

Function 20, 6-93

Function 21, 6-113

Function 22, 6-79

Function 23, 6-98

Function 24, 6-56

Function 25, 6-54

Function 26, 6-73

Function 27, 6-46

Function 28, 6-62

Function 29, 6-58

Function 30, 6-65

Function 31, 6-48

Function 32, 6-112

Function 33, 6-95

Function 34, 6-116

Function 35, 6-102

Function 36, 6-92

Function 37, 6-57

Function 38, 6-45

Function 39, 6-53

Function 40, 6-121

Function 42, 6-76

Function 43, 6-109

Function 44, 6-82

Function 45, 6-75

Function 46, 6-63

Function 47, 6-141

Function 48, 6-52

Function 50, 6-175

Function 51, 6-74

Function 52, 6-72

Function 53, 6-138

Function 54, 6-134

Function 55, 6-136

Function 56, 6-133

Function 57, 6-137

Function 58, 6-135

Function 59, 6-156

Function 99, 6-107

Function 100, 6-60

Function 101, 6-55

Function 102, 6-105

Function 103, 6-119

Function 104, 6-188

Function 105, 6-185

Function 106, 6-91

Function 107, 6-177

Function 109, 6-28

Function 110, 6-25

Function 111, 6-39

Function 112, 6-128

Function 128, 6-131

Function 129, 6-131

Function 130, 6-132

Function 131, 6-41

Function 132, 6-43

Function 133, 6-42

Function 134, 6-168

Function 135, 6-171

Function 136, 6-167

Function 137, 6-172

Function 138, 6-165
 Function 139, 6-173
 Function 140, 6-166
 Function 141, 6-154
 Function 142, 6-155
 Function 143, 6-161
 Function 144, 6-145
 Function 145, 6-158
 Function 146, 6-23
 Function 147, 6-26
 Function 148, 6-36
 Function 149, 6-21
 Function 150, 6-142
 Function 151, 6-159
 Function 152, 6-86
 Function 153, 6-27
 Function 154, 6-178
 Function 155, 6-187
 Function 156, 6-157
 Function 157, 6-139
 Function 158, 6-122
 Function 159, 6-124
 Function 160, 6-126
 Function 161, 6-123
 Function 162, 6-24
 Function 163, 6-176, 6-182
 Function 164, 6-125
 F[^] interface attribute, 6-76
 F_, 2-2
 F_ATTRIB system call, 2-14, 2-30,
 2-31, 2-38, 6-65, 6-83, 6-98
 F_CLOSE system call, 2-30, 2-33,
 2-39, 6-68
 F_DELETE system call, 2-30, 6-70,
 6-80
 F_DMAGET system call, 6-73
 F_DMAOFF, 6-156
 F_DMAOFF system call, 5-6, 6-74,
 6-75
 F_DMASEG, 6-73, 6-156
 F_DMASEG system call, 5-6, 6-74
 F_ERRMODE system call, 2-29, 2-49,
 6-45, 6-75
 F_FLUSH system call, 2-39
 F_LOCK, 2-35
 F_LOCK system call, 2-26, 2-34, 2-36,
 6-76, 6-82
 F_MAKE, 6-76
 F_MAKE system call, 2-10, 2-14,
 2-21, 2-22, 2-27, 2-38, 6-79, 6-93,
 6-113
 F_MULTISEC system call, 2-34, 6-82,
 6-93, 6-95, 6-113
 F_OPEN, 6-76
 F_OPEN call, 2-26
 F_OPEN system call, 2-9, 2-10, 2-14,
 2-26, 2-27, 2-31, 2-38, 6-66, 6-83,
 6-93, 6-109, 6-113, 6-143
 F_PARSE system call, 2-6, 3-1, 6-87,
 6-143
 F_PASSWD, 6-98
 F_PASSWD system call, 2-23, 6-61,
 6-65, 6-91, 6-107
 F_RANDREC system call, 6-92
 F_READ system call, 2-34, 6-93
 F_READRAND system call, 2-34,
 6-96
 F_RENAME system call, 2-12, 2-30,
 2-31, 6-98
 F_SFIRST system call, 2-14, 2-15,
 2-20, 2-23, 2-25, 2-38, 6-66, 6-70,
 6-100
 F_SIZE system call, 6-102
 F_SNEXT system call, 2-14, 2-15,
 2-20, 2-23, 2-25, 2-38, 6-66, 6-70,
 6-100, 6-104
 F_TIMEDATE system call, 2-25,
 6-105
 F_TRUNCATE system call, 2-30,
 6-107
 F_UNLOCK, 2-35

F_UNLOCK system call, 2-26, 2-34,
2-35, 2-36, 2-37, 6-84, 6-109
F_USERNUM system call, 2-17,
6-112
F_WRITE system call, 2-34, 6-113
F_WRITERAND system call, 2-34,
6-94, 6-102, 6-116
F_WRITEXFCB system call, 2-21,
2-22, 6-119
F_WRITEZF system call, 2-34, 6-94,
6-121

G

G_Form, 3-3
G_Type field, 3-2
GENCCPM, 2-29, 3-1, 5-1, 5-3, 5-11,
6-181
GENCMD, 4-6, 4-9, 5-2
generic category, 2-7
Group Descriptor, 3-3
Group Descriptor Format
(Figure 3-2), 3-3
G_Length, 3-4
G_Max, 3-4
G_Min, 3-4

H

H86 filetype, 4-6
Hard Disk, 6-51
hardware initialization, 6-180
Header Record, 3-3
CMD file, 4-1, 4-7
header
RSP, 5-2
HEX file, 4-6, 4-7
highest priority process, 1-6
hour of day, 6-186

HOUR
TOD field, 6-186

I

Illegal ? in FCB error, 2-46
independent group, 3-7
initial flags, 6-152
initial stack area, 4-2
initial stack
8080 model, 4-2
initial values
instruction pointer, 4-1
segment registers, 4-1
stack pointer, 4-1
initialization
hardware, 6-180
initialize directory, 2-39
Instruction Pointer, 4-3, 6-153
INT 0, 6-153
INT 1, 6-153
INT 3, 6-153
INT 4, 6-153
INT 224, 1-12, 6-153
INT 225, 6-153
Intel hexadecimal file format, 4-9
Intel utilities, 4-7
Intel
small model, 4-1
interface attribute
F5', 6-68, 6-70, 6-83
F6', 6-70, 6-83
F7', 6-84
F8', 6-84
interface attributes, 2-14, 2-16, 2-27,
6-65
Interrupt Return instruction, 6-152,
6-180
interrupt returns, 6-180
interrupt vectors, 6-153

- interrupt
 - logical, 1-2
 - physical, 1-2
 - types, 6-43
- interrupts enabled, 5-9
- Invalid Drive error, 2-44
- IO_CONIN
 - XIOS, 1-10
- IP, 6-153
- IP flag, 6-42
- IP register, 4-3
 - Small Model, 4-4
- IP
 - instruction pointer, 6-152
- IRET instruction, 5-9
- IRET structure, 5-11

J

- JMPF PDISP instruction, 6-180

K

- KEEP Flag, 5-8
- KEEP flag, 5-9, 6-149, 6-161
- KERNEL flag, 6-149

L

- label
 - directory, 2-18
- last record byte count, 6-65
- last record number, 6-107
- LCB, 1-10
 - SYSDAT field, 6-183
- Ld Addr, 4-11
- Least Recently Used order, 3-2

LENGTH

- MCB, 6-129
- line feed, 2-9, 6-32, 6-33, 6-34, 6-90
- line-editing, 6-33, 6-34
- LINK field, 6-146, 6-182
 - RSP header, 5-7, 5-11
- Link list root, 6-181
- Link list
 - memory allocation units, 6-182
- LINK
 - PD field, 6-147
- LIST, 6-175
- List Control Block, 1-10, 6-181, 6-183
- list device, 1-11, 6-122, 6-123, 6-124, 6-126, 6-127
- List Device I/O System Calls, 6-122
- List Device System Calls, 6-2, 6-8
- list devices
 - number of XIOS, 6-182
- List field
 - process descriptor, 1-11
- list I/O, 1-10
- LIST
 - PD field, 6-150
- lock existing records only, 6-76
- Lock List, 2-27, 2-28, 2-29, 2-30, 2-33, 2-37, 2-41, 2-42, 6-45, 6-53, 6-77, 6-81, 6-85, 6-110, 6-161, 6-181
- lock logical records, 6-76
- lock operations, 2-36, 2-37
- Locked, 2-2
- Locked mode, 2-26, 2-30, 6-19, 6-80, 6-83
- locked records
 - maximum number, 6-183
- locks
 - exclusive, 2-35
 - shared, 2-35
- LOCK_MAX
 - SYSDAT field, 6-183
- log-in drive, 2-3

- log-in operation, 2-39
- logged-in, 2-39
- logical console, 6-37, 6-39
- logical drives, 2-8
- logical extent, 6-113
- logical interrupt, 1-2, 6-42
- logical list device, 6-128
- logical message, 6-169
- logical record size, 2-37
- Login Vector, 6-44, 6-56
- lowercase, 2-6, 2-7
- LRU, 3-2
- LST:, 6-128
- LUL
 - SYSDAT field, 6-181
- L_ATTACH, 6-127
- L_ATTACH system call, 6-122
- L_CATTACH system call, 6-123
- L_DETACH system call, 6-124
- L_GET system call, 6-125
- L_SET system call, 6-126
- L_WRITE system call, 6-127
- L_WRITEBLK system call, 6-128

M

- M value, 4-8
- M80 byte, 3-7
- machine code
 - Small Model, 4-4
- make system queue, 6-168
- MAL
 - SYSDAT field, 6-182
- MATCH
 - C_ASSIGN system call, 6-22
- MAX number of paragraphs, 6-131
- MAX
 - C_READSTR, 6-34
 - MPB, 6-130

- MCB—Memory Control Block
 - (Figure 6-7), 6-128
- MC_ABSALLOC system call, 6-133
- MC_ABSMAX system call, 6-134
- MC_ALLFREE system call, 6-135
- MC_ALLOC system call, 6-136
- MC_FREE system call, 6-137
- MC_MAX system call, 6-138
- MDUL
 - SYSDAT field, 6-182
- media change, 2-3, 2-29, 2-39, 2-40, 2-41, 2-42
- media
 - nonremovable, 6-50
- MEM, 1-4, 1-9
- MEM field
 - Process Descriptor, 5-4
- MEM
 - DP field, 6-149
- memory, 3-7
- memory allocation, 1-13
- Memory Allocation System Calls
 - MP/M Compatible, 6-9
 - CP/M Compatible, 6-9
- memory allocation units, 6-182
- Memory Control Block, 6-128
 - Definition, 6-129
- Memory Descriptors
 - unused, 6-182
- Memory Management System Calls,
 - 6-2, 6-3
- Memory Management Module, 1-4
- memory model, 4-1
 - RSP, 5-1
- Memory Module, 1-9
- Memory Parameter Block Definition,
 - 6-130
- memory partitions
 - free, 6-182
- memory protection, 6-146
- Memory Segment Descriptors, 6-149

Memory System Calls, 6-128
 memory
 absolute, 6-134
 initialization, 3-1
 largest available region, 6-138
 maximum per process, 6-181
 message
 length, 5-10, 6-169
 maximum number, 6-169
 zero-length, 1-8
 MFL
 SYSDAT field, 6-182
 MFPB—M_FREE Parameter Block
 (Figure 6-9), 6-132
 Mhhh parameter, 4-7
 MIN length, 6-131
 MIN
 MPB, 6-130
 TOD field, 6-185
 minimum memory value, 4-8
 minimum memory requirement, 4-7
 minute of hour, 6-183, 6-185
 MMP
 SYSDAT field, 6-181
 modes
 file open, 2-26
 MP/M Compatible Memory
 Allocation System Calls, 6-9
 MP/M-86 memory allocation scheme,
 6-128
 MPB—Memory Parameter Block
 (Figure 6-8), 6-129
 MSGLEN
 QD field, 6-169
 multi-user, 1-1
 multiple programs, 1-2
 Multisector count, 2-3, 2-34, 2-35,
 2-36, 6-12, 6-13, 6-73, 6-76, 6-82,
 6-93, 6-117, 6-118
 Multisector I/O, 2-34
 mutual exclusion queues, 1-7, 1-8

 MX queue, 1-8
 MXdisk, 1-8
 M_ALLOC system call, 6-131
 M_FREE system call, 6-131, 6-132

N

 NAME field, 5-8
 directory label, 2-19
 APB field, 6-140
 CPB field, 6-160
 C_ATTACH, 6-23
 DP field, 6-149
 PD, 5-3
 QD field, 6-169
 QPB field, 6-163
 queue, 6-169
 RSP PD, 5-8
 NCCB
 SYSDAT field, 6-181
 NCHAR
 C_READSTR, 6-34
 NCIODEV
 SYSDAT field, 6-183
 NCONDEV
 SYSDAT field, 6-183
 NCP byte
 field, 5-3
 RSP header, 5-3
 networking interfaces, 1-5
 NFLAGS
 SYSDAT field, 6-181
 NLCB
 SYSDAT field, 6-181
 NLSTDEV
 SYSDAT field, 6-183
 NMSGs
 QD field, 6-169
 no data, 6-94

No Room In System Lock List error, 2-46
non-8080 model, 3-7
noninterrupt-driven devices, 6-41
Nonremovable Media Drives, 6-50
null character, 6-90
NVCNS
 SYSDAT field, 6-181
NVCNS field, 5-3

O

OFF, 6-50
OFF_8087
 SYSDAT, 6-184
OH86 utility, 4-9
one second flag
 Flag 2, 1-8
open disk files
 maximum number, 6-183
open file, 2-2
Open File Drive Vector, 6-183
Open File Limit Exceeded error, 2-46
open mode, 2-2, 2-26
open verification, 2-29
OPEN_FILE
 SYSDAT field, 6-183
OPEN_MAX
 SYSDAT field, 6-183
Operating System Version Number
 Format (Figure 6-19), 6-176
OS type, 6-174, 6-176
os version, 6-176
Output Delimiter, 6-25
owner
 queue message, 1-8
OWNER_8087
 SYSDAT, 6-183

P

P1 Len, 3-7
P2 Len, 3-7
PARAM field
 CPB, 6-160
PARAM
 CPB field, 6-160
parameter passing, 6-140
PARENT
 PD field, 6-149
parent/child relationship, 3-8
parentheses, 2-6
parse file specification, 2-3
Parse Filename Control Block, 6-86
partial close, 2-30, 2-33, 6-68
password, 2-1, 2-2, 3-7, 6-61, 6-65, 6-78, 6-98
 default, 2-3, 2-23
 length, 3-7
 mode, 6-79, 6-105
password error, 2-45
password field, 2-5, 6-89
 directory label, 2-19
Password field
 XFCB, 2-21
password protection, 1-10, 2-3, 2-22, 6-80
password support, 2-18
PD, 1-5, 5-1
PD—Process Descriptor
 (Figure 6-12), 6-146
PD address, 6-157
PD table, 6-145, 6-149, 6-161
PD
 APB field, 6-140
 C_ASSIGN, 6-22
PDADDRESS
 RSP Command Queue Message, 5-5

PDISP
 SYSDAT field, 6-180
 permanent drive, 2-39, 2-40, 2-42
PFCB—Parse Filename Control
 Block (Figure 6-6), 6-86
 Physical and Extended Errors, 2-49
 physical error, 2-43, 2-49, 2-50
 Physical Input Process, 1-10, 6-31
 physical interrupt, 1-2
 Physical Record Mask, 6-50
 Physical Record Shift Factor, 6-50
 physical records, 2-38
 PIN, 1-10, 1-11, 6-31
 PIP utility, 2-15, 2-34
PLR
 SYSDAT field, 6-182
PM field
 XFCB, 2-21
 Poll List, 6-147
POLL
 Activity code, 6-147
 List Root, 6-182
 printer, 1-11,
 echo, 6-29
 priority
 highest, 6-158
 lowest, 6-158
 transient process, 5-4, 6-158
PRIORITY field, 5-8
PRM, 6-49, 6-51
 process, 1-2, 2-28, 2-35
 Process Descriptor, 1-5, 5-1, 6-144,
 6-145, 6-146, 6-161, 6-178
 address, 1-8, 6-140, 6-157
 Process Descriptor
 initialization, 3-1
 unused, 6-182
 Process ID
 C_ASSIGN, 6-22
 Process Keep flag, 1-11
 process name, 6-149
 aborted, 1-8
 priority, 6-154
 privileged, 5-10
 register values, 6-152
 resources, 6-161
 scheduling, 6-148
 Process/Program System Calls, 6-3,
 6-11
 program, 1-2
 Program Flag
 CMD header record, 3-2
PSH, 6-49, 6-51
PUL
 SYSDAT field, 6-182
P_ABORT, 1-11
P_ABORT system call, 6-140
P_CHAIN system call, 2-17, 6-141
P_CLI system call, 1-5, 2-6, 2-7, 2-17,
 2-32, 3-1, 4-2, 4-3, 4-4, 4-5, 5-4,
 5-5, 5-6, 6-32, 6-73, 6-82, 6-143,
 6-144
P_CREATE, 6-145
P_CREATE system call, 3-1, 5-1, 5-4,
 5-8, 5-10, 6-146, 6-149, 6-157
P_DELAY system call, 1-9, 6-154
P_DISPATCH system call, 6-155
P_LOAD system call, 1-5, 3-5, 4-2,
 4-6, 6-143, 6-156
P_PDADR system call, 5-5, 6-157
P_PRIORITY system call, 5-8, 6-158
P_RPL system call, 6-160
P_TERM, 3-1, 4-2, 6-162
P_TERM system call, 6-32, 6-140,
 6-141, 6-161
P_TERMCPM, 4-2
P_TERMCPM system call, 6-162
P_TERMCPM
 CP/M-86, 6-162

Q

QD—Queue Descriptor (Figure 6-16),
6-168

QLR
SYSDAT field, 6-182

QMAU
SYSDAT field, 6-182

QPB, 6-171

QPB—Queue Parameter Block
(Figure 6-15), 6-163

qualified reset, 2-40

question mark, 2-6

queue buffer, 1-7, 6-145, 6-169

queue descriptor, 1-7, 1-8; 6-147,
6-168

unused, 6-182

queue flags, 6-169

ID Field, 6-171

List Root, 6-182

Management, 1-7

Management System Calls, 6-3

message, 1-6, 1-7

Message Buffer, 6-163

name, 1-7, 6-163, 6-169

Parameter Block, 5-10, 6-163

System Calls, 6-12, 6-163

QUEUID
QPB field, 6-163

QUL
SYSDAT field, 6-182

Q_CREAD system call, 5-5, 6-165

Q_CWRITE system call, 5-5, 6-166

Q_DELETE system call, 5-9, 6-167

Q_MAKE system call, 1-7, 5-10, 6-168

Q_OPEN, 5-5, 6-163

Q_OPEN call, 6-172, 6-173

Q_OPEN system call, 6-160, 6-170,
6-171

Q_READ, 1-6

Q_READ system call, 5-5, 6-165,
6-172

unconditional, 1-8

Q_WRITE, 1-6

Q_WRITE system call, 5-5, 6-166,
6-173

R

R/O drive test, 2-42

R/O Vector, 6-58

R0

R1 field, File ID, 6-80

R1,R2 field, 6-18

R1,R2 field, FCB, 2-12

R1,R2 fields, 6-92

random, 2-2

read, 2-9, 6-12

Random Record Field, 2-36

FCB, 2-35

Random Record Number, 2-9, 2-37,
3-8, 6-76, 6-92, 6-96, 6-102, 6-109,
6-111, 6-117

FCB, 2-12

raw console output, 6-29

mode, 6-31

RC field

FCB, 2-11

XFCB, 2-21

read message, 6-172

read mode, 2-22, 6-80, 6-105

Read Queue List, 6-147

read record, 2-2, 6-93

Read-Only, 2-2, 2-40, 6-65
 mode, 2-26
 attribute, 2-15, 2-26
 attribute TI', 6-84
 attribute TI', 2-15
 drive, 6-62
 file, 2-11, 6-76
 mode, 2-35, 6-83
 Vector, 6-44
 Read-Write, 2-40
 Read-Write, 6-47
 Read-Write state, 6-62
 Read/Only Disk error, 2-44
 File error, 2-44
 Ready List, 1-5, 1-6, 1-7, 1-9, 6-147
 Ready List Root, 6-182
 ready process, 1-5
 Real-time Monitor, 1-4, 1-5
 real-time process control, 1-2
 window, 1-13
 Rec Len, 4-11
 Rec Mark, 4-11
 Rec Type, 4-11
 record blocking, 2-38, 6-82
 record count
 file, 2-9
 first, 2-9
 locking, 2-28, 2-36
 physical, 2-38
 size, 2-2, 2-37
 unlocking, 2-36
 REDRAW, 6-35
 reentrant, 6-149, 6-160
 reentrant RSP, 5-4
 register AL, 2-47
 register contents preserved, 1-13
 register initialization, 5-8, 5-9
 removable drive, 2-40, 2-42
 reset
 drive, 2-39
 Resident Procedure Library, 6-160
 resident system process, 1-2, 3-1, 5-1,
 6-143
 resources
 process, 6-161
 RESTRICTED flag, 5-10
 RETF instruction, 4-2, 6-180
 RETURN, 6-35
 Return and Display Error mode, 2-43
 Return and Display mode
 BDOS Error mode, 6-75
 return codes, 2-47
 Return Error mode, 2-43, 2-49
 BDOS Error mode, 6-75
 Revision Level, 6-176
 RLR
 SYSDAT field, 6-182
 roots of system lists, 6-178
 round-robin scheduling, 6-155
 RPL, 6-160
 RS field
 FCB, 2-11
 RSP, 1-2, 6-143
 bit, 5-9
 CMD Header Record, 5-2
 ECHO, 5-1
 first, 6-181
 multiple copies, 5-3
 shared code, 5-4
 8080 Model, 5-2, 5-3
 Small Model, 5-2, 5-4
 RSP Command Queue, 5-4, 5-5, 5-6,
 5-9

- RSP Command Queue Message
 - (Figure 5-3), 5-5
 - Data Segment (Figure 5-4), 5-7, 6-180
 - Flag, 5-5
 - header, 5-2, 5-3, 5-6, 5-7
 - Header Format (Figure 5-2), 5-3
 - memory models, 5-1
 - Process Descriptor, 5-4, 5-8
 - queue, 6-143
 - stack, 5-9
 - type, 3-1
 - UDA, 5-6, 5-7
- RSPSEG field, 5-11
- RSPSEG
 - SYSDAT field, 6-180
- RTM, 1-4, 1-5, 1-8
- RUB/DEL, 6-34
- RUN state, 6-41
- RUN
 - Activity code, 6-147
- running process, 1-1, 1-5

- S**

- SI
 - S2 fields, directory label, 2-19
 - S2 fields, XFCB, 2-21
- screen switch, 1-10, 1-11
- SDATVAR field
 - RSP header, 5-3
- SEC
 - TOD field, 6-186
- second flag, 1-8
- second of minute, 6-183, 6-186
- seconds, 6-187
- Sectors Per Track, 6-49
- security
 - file, 2-27
- segment addresses, 6-153

- Segment Base Address, 4-9
- segment register initialization, 4-2
- SEG_8087
 - SYSDAT, 6-184
- sequential, 2-2
 - access, 6-12
 - I/O processing, 2-34
 - read, 2-9
 - write, 6-79
- serial number, 6-177
- SERIAL Number Format
 - (Figure 6-20), 6-177
- SET command, 2-23
- SET utility, 2-32, 2-33
- SFCB, 2-18, 2-24, 6-17
- SFCB Subfields (Figure 2-5), 2-24
- SFCBs, 6-61
- shared code, 1-2, 3-2
 - file access, 1-10
 - file system, 1-2
 - List, 3-2
 - RSPs, 6-149
 - locks, 2-35, 6-77
- SI
 - UDA field, 6-152
- single-user, 1-1
- size
 - physical records, 2-38
 - record, 2-2, 2-37
- Small Memory Model, 3-5, 4-4
- Small Model, 1-12, 4-2
- source files, 2-9
- SP field
 - UDA, 5-9, 6-152
- sparse file, 2-9
- SPT, 6-49
- SS and SP registers
 - Small Model, 4-4
 - UDA field, 6-153
- stack area, 6-144
- stack pointer, 6-152

- Stack Segment, 5-1, 6-152
- stack
 - RSP, 5-9
- start address, 4-7, 4-9
- START field, 6-131
- START paragraph, 6-132
 - MPB, 6-130
- STAT
 - PD field, 6-147
- state
 - reset, 2-39
- Status Word
 - UDA 8087 extension, 6-153
- string delimiter, 6-40
- SUP, 1-4, 1-5
- SUP ENTRY
 - SYSDAT field, 6-180
- Supervisor, 1-4, 1-5
 - Code Segment, 6-180
 - entry point, 6-180
- suspended process, 1-5
- SW
 - UDA 8087 extension, 6-153
- switch screen, 1-11
- synchronization, 1-2
- SYS Flag, 5-8
- SYS flag, 6-140, 6-148
- SYSDAT Table (Figure 6-21), 6-179
- SYSDAT, 2-25, 5-11
 - H(elp) option, 1-14
 - M(emory) option, 1-14, 1-15
 - SYSDAT field, 6-181
- SYSDISK
- SYSTAT, 1-14
 - O(overview) option, 1-15
 - P(rocess) option, 1-15
 - Q(ueues) option, 1-15
 - U(ser Procèsses) option, 1-15
- System, 6-65
- system attribute, 2-15
- SYSTEM attribute, 6-143
- system attribute t2', 6-84
- system calls 3, 6-1, 6-18, 6-21
 - conventions, 1-12
- system call register initialization, 1-13
- System Call summary, 6-14
- System Data Area, 5-7, 5-10
- System Data Segment, 5-11, 6-145, 6-170
 - address, 6-178
- system disk, 6-143
 - default, 6-181
- System file, 2-11
 - user-zero, 2-15
- SYSTEM flag, 6-161
- system
 - flags, 6-181
 - generation, 5-1
 - lists, 1-5, 1-6
 - process, 6-148
 - processes, 1-2
 - queue, 1-2, 1-13, 6-170
 - Status, 1-14
 - System Calls, 6-3, 6-13
 - ticks, 6-162, 6-155
 - ticks per second, 6-181
 - time and date, 6-185
 - timing, 1-8, 1-9
 - tracks, 2-8
- S_BDOSVER, 6-182
- S_BDOSVER system call, 6-174
- S_BIOS system call, 6-175
- S_OSVER, 6-182
- S_OSVER system call, 6-176
- S_SERIAL system call, 6-177
- S_SYSDAT system call, 5-7, 6-178

T

T1', 2-15

T1' attribute, 2-26

T1'-T3', 2-14, 6-65

FCB, 2-11

T2', 2-15, 2-18

T3', 2-15

TAB, 6-35, 6-90

characters, 6-32

expansion, 6-29, 6-38, 6-39

TABLE flag, 6-149

TEMP DISK

SYSDAT field, 6-181

TERM

APB field, 6-139

Terminal Message Processes, 1-11

Terminal Message Processor, 1-4, 3-1

termination

character, 6-32, 6-33

code, 6-139, 6-161, 6-162

THRDRT

SYSDAT field, 6-182

THREAD

field, 6-182

list, 6-22, 6-139, 6-147

List Root, 6-182

PD field, 6-147

tick flag, 1-9

Tick Interrupt Handler

XIOS, 1-8, 1-9

TICKS/SEC

SYSDAT field, 6-181

time and date, 1-2, 1-8, 6-105, 6-185,

6-187

time of day, 1-8

time stamp

directory label, 2-25

Time System Calls, 6-3, 6-13

timing functions, 1-2

TMP, 1-4, 1-11, 2-17, 3-1, 5-5, 5-9,
6-112, 6-161

priority, 6-144

RSP, 5-3

TOD—Time-of-Day Structure

(Figure 6-22), 6-185

TOD_DAY

SYSDAT field, 6-183

TOD_HR

SYSDAT field, 6-183

TOD_MIN

SYSDAT field, 6-183

TOD_SEC

SYSDAT field, 6-183

TPA, 6-145

Track Offset, 6-51

Transient Execution Models, 4-1

Process Area, 6-145

processes, 1-2, 1-5

program, 1-12, 3-1

truncate file, 2-1, 2-2

TS1 field

directory label, 2-19

TS2 field

directory label, 2-19

type field

directory label, 2-19

XFCB, 2-21

TYPE utility, 2-9

T_GET system call, 2-25, 6-186

T_SECONDS system call, 6-187

T_SET system call, 6-188

U

UDA, 1-5, 1-6, 5-1, 6-135, 6-144,
6-145

UDA—User Data Area (Figure 6-23),
6-151

UDA SEGMENT field, 5-8

UDA
8087, 3-1, 3-2
initialization, 3-1
PD field, 6-149
RSPs, 3-1
unallocated data block, 6-121
unconditional read
queue, 1-7
unlock operations, 2-36
records, 6-111
unlocked, 2-2
mode, 1-10, 2-12, 2-26, 2-35, 2-37,
6-79, 6-82
unused Process Descriptors, 6-182
unused Queue Descriptors, 6-182
unused Memory Descriptors, 6-182
unwritten random records, 6-121
update date and time stamp, 6-17,
6-114
update stamp, 6-80
field, 2-19
time stamp, 2-24
Upper Segment Base Address, 4-12
USBA, 4-12
User 0, 2-18, 6-83
user attributes, 2-15
User Data Area, 1-5, 3-1, 5-1, 6-135,
6-145, 6-149, 6-151
RSP, 5-9
user default disk, 6-181
directories, 2-17
number, 1-11, 2-1, 2-3, 6-82
number conventions, 2-17
terminal, 1-1
zero, 6-82
user processes priorities, 6-148
User System Stack, 6-152
USER SYSTEM STACK
UDA field, 6-153
USER
PD field, 6-149

user-zero system files, 2-15

V

VERNUM
SYSDAT field, 6-182
version number, 6-174, 6-182
version string address, 6-182
version
os, 6-176
VERSION
SYSDAT field, 6-182
VINQ, 1-11
virtual console, 1-1, 1-2, 6-175, 6-181
Virtual Console Input Queue, 1-11
Virtual Console Screen Management,
1-10
Virtual Console Screen Manager, 1-4
virtual environments, 1-1
virtual file size, 6-18
Virtual OUTput processes, 1-10
VOUT, 1-10

W

wildcard file specifications, 6-70
window
real-time, 1-13
write data records, 6-113
write, 6-173
write mode, 2-22, 6-80, 6-105
Queue List, 6-147
record, 2-2
sequential, 6-79
zeroes, 6-121

X

- X value, 4-8
- XFCB, 2-18, 2-20, 6-79, 6-81
 - Extended File Control Block (Figure 2-3), 2-20
 - Create or access time stamp field, 6-105
 - password mode, 6-119
 - Update time stamp field, 6-105
- Xhhhh parameter, 4-7
- XIOS, 1-4, 1-10, 1-11, 6-41, 6-43, 6-175
 - ALLOC buffers, 6-181
 - ENTRY, 6-180
 - Header, 6-181
- XIOS INIT, 6-180
- XIOS Initialization entry point, 6-180
- XPCNS
 - SYSDAT, 6-184

Z

- Zeroes, 4-11

Reader Comment Card

We welcome your comments and suggestions. They help us provide you with better product documentation.

Date _____


1. What sections of this manual are especially helpful?

2. What suggestions do you have for improving this manual? What information is missing or incomplete? Where are examples needed?


3. Did you find errors in this manual? (Specify section and page number.)

Concurrent CP/M™ Operating System Programmer's Reference Guide
First Edition: January, 1984
1034-2023

From: _____



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS / PERMIT NO. 182 / PACIFIC GROVE, CA

POSTAGE WILL BE PAID BY ADDRESSEE

 **DIGITAL RESEARCH™**

P.O. Box 579
Pacific Grove, California
93950

Attn: Publications Production

