Systems

A Guide to the IBM System/370 Model 165

IBM

Systems

A Guide to the IBM System/370 Model 165

This guide presents hardware, programming systems, and other pertinent information about the IBM System/370 Model 165 that describes its significant new features and advantages. Its contents are intended to acquaint the reader with the Model 165 and to be of benefit in planning for its installation.

IBM

First Edition (June 1970)

This guide is intended for planning purposes only. It will be updated from time to time to reflect system changes; however, the reader should remember that the authoritative sources of system information are the Systems Reference Library (SRL) publications for the Model 165, its associated components and its programming support. These publications will first reflect such changes.

Copies of this and other IBM publications can be obtained through IBM branch offices.

A form has been provided at the back of this publication for readers' comments. If this form has been removed, address comments to: IBM Corporation, Technical Publications Department, 112 East Post Road, White Plains, New York 10601.

PREFACE

It is assumed that the reader of this document is familiar with System/360. The reader should have a general knowledge of System/360 architecture, channels, I/O devices, and programming systems support. This guide highlights only those Model 165 hardware, I/O, and programming systems features that are different from those of System/360 models and discusses their significance. Additional, more detailed information regarding System/370 Model 165 hardware and programming systems support can be found in the following SRL publications:

IBM System/370 Model 165 Functional Characteristics (GA22-6935)

IBM System/370 Principles of Operation (GA22-7000)

IBM System/370 System Summary (GA22-7001)

IBM System/370 I/O Configurator (GA22-7002)

Component Summary: 3830 Storage Control, 3330 Disk Storage (GA26-1592)

IBM System/360 Component Description: 2835 Storage Control and 2305 Fixed Head Storage Module (GA26-1589)

3211 Printer and 3811 Control Unit Component Description (GA24-3543)

IBM Component Description: 3803/3420 Magnetic Tape Subsystem (GA32-0020)

Form-Design Considerations - System Printers (GA24-3488)

Emulating the 7070/7074 on the IBM System/370 Model 165 using OS/360 (GC27-6948)

Emulating the 709, 7090, 7094, 7094II on the IBM System/370 Model 165 using OS/360 (GC27-6951)

Emulating the 7080 on the IBM System/370 Model 165 using OS/360 (GC27-6952)

IBM System/360 Operating System:

Planning for the IBM 3211 Printer, Data Management Macro Instructions and Services (GC21-5008)

CONTENTS

FIGURES

TABLES

SECTION 01: SYSTEM HIGHLIGHTS

The System/370 Model 165 is designed to enhance, extend, and broaden the successful concepts of System/360 architecture. It is a high-speed, large-scale growth system for System/360 Model 65 and 75 users that provides significant price performance improvement without the necessity of major reprogramming. The Model 165 is a general purpose system and offers high performance for both commercial and scientific applications. It is compatible with the System/370 Model 155.

Transition from System/360 models to the System/370 Model 165 can be accomplished with a minimum of effort because most current System/360 user programs, I/O devices, and programming systems are upward compatible with the new system. Upward transition from a Model 165 to a Model 195 can also be accomplished. Additional capabilities will be added to OS to support new features of the Model 165, thereby providing proven operating system performance as well as continuity.

Transition with little or no reprogramming is also provided for 705I/II/III/7080, 7070/7074, and 709/7090/7094/7094II users who are presently emulating on System/360. Emulators that operate under OS control on the Model 165 are provided for these systems.

Highlights of the Model 165 are as follows.

- Upward compatibility with most System/360 architecture and programming has been maintained.
- Internal performance is approximately two to five times that of the Model 65.
- CPU features of the System/370 Model 165 are as follows:

  The Model 165 standard instruction set includes new general purpose instructions in addition to the powerful System/360 instruction set. These instructions enhance decimal arithmetic performance, eliminate the need for multiple move or compare instructions or move subroutines, and facilitate record blocking and deblocking, field padding, and storage clearing.

  Extended precision floating point is a standard feature to provide precision of up to 28 hexadecimal digits, equal to approximately 34 decimal digits.

  A high-speed multiply feature is available to provide faster execution of binary and floating-point arithmetic operations. Execution speed increases by a factor of 2 to 3.

  An interval timer of 3.33 ms resolution to improve job accounting accuracy is standard. A 16.6 ms resolution timer is standard on the Model 65.

  A time of day clock is included to provide more accurate time of day values than the interval timer. It has a 1 microsecond resolution.

  Separate instruction and execution units are implemented that provide overlap of instruction fetching, instruction decoding, operand fetching, and instruction execution to increase internal performance.

CPU retry of most failing CPU hardware operations is handled automatically by the hardware without programming assistance.

Writable control storage (WCS) is included in addition to read-only storage (ROS) to contain new Model 165 instructions, emulator microcode, and CPU diagnostics.

- Relocatable emulators are provided that operate under OS control. Concurrent execution of System/370 programs with 7000-series programs is supported. A 7080, a 7070/7074, and a 709/7090/7094/7094II emulator are available on a mutually exclusive basis.

- A free-standing 3066 System Console is required. Its features are:

  A buffered cathode ray tube and an alphameric keyboard for rapid operator/system communication

  An indicator viewer to display system status

  A system activity monitor to provide, via the system activity meter, visual display of average system activity

  A microfiche document viewer for CE use

  A processor storage configuration plugboard

  A device for loading WCS and diagnostic routines

- Channel features of the Model 165 are as follows:

  2870 Multiplexer Channels, 2860 Selector Channels, and the new 2880 Block Multiplexer Channels can be attached - for a total of seven addressable channels. A single 2880 channel can operate at a 3 MB rate with attachment of an optional feature.

  The Extended Channels feature permits a Model 165 to have up to twelve addressable channels, which can support an aggregate channel data rate in excess of nine megabytes per second.

  The 2880 Block Multiplexer Channel is a superset of the 2860 Selector Channel. When used in conjunction with rotational position sensing devices, it can increase total system throughput by permitting more data to enter and leave the system in a given time period than can the 2860. A single 2880 channel can support interleaved, concurrent execution of multiple high-speed I/O operations.

  Channel retry data is provided after channel errors so that error recovery routines can retry I/O operations.

- Storage features offered by the Model 165 are as follows:

  A two-level memory system, consisting of fast, large-size processor (main) storage used as backing storage for a smaller, very high-speed buffer storage, is implemented. The CPU works mostly with the buffer so that the effective access time for data is reduced to a fraction of the processor storage cycle time.

  8K or 16K bytes of 80-nanosecond monolithic buffer storage is available (8K is standard). The CPU can initiate a request for eight bytes from the buffer every 80 nanoseconds.

512K to 3072K of four-way, doubleword interleaved, two-microsecond processor storage is available - three times the maximum available on the Model 65.

Byte boundary alignment is permitted for the operands of non-privileged instructions to eliminate the necessity of adding padding bytes within records or to blocked records for the purpose of aligning fixed or floating-point data.

Error checking and correction (ECC) hardware that automatically corrects all single-bit processor storage errors and detects all double-bit and most multiple-bit errors is standard.

- I/O devices for the Model 165 are as follows:

Most currently announced I/O devices for System/360 Models 65 and above can be attached.

The new 3330 facility is available for attachment to 2880 channels. It offers significantly faster seeks and more than twice the data rate of the 2314 facility, more than three times the capacity of the 2314, and automatic error correction features. The new rotational position sensing and multiple requesting capabilities announced for the 2305 facility are standard.

The 3330 has an 806 KB data transfer rate, average seek time of 30 ms, and full rotation time of 16.7 ms. Up to 800 million bytes can be contained on an eight drive facility.

The 2305 facility Models 1 and 2 can be connected to 2880 channels to provide significantly faster data transfer operations.

The Model 1 has a 3 megabyte data rate, a maximum module capacity of 5.4 million bytes, and average access time of 2.5 ms.

The Model 2 has a 1.5 megabyte data rate, a maximum module capacity of 11.2 million bytes, and average access time of 5 ms.

The new high-speed 3211 Printer with a tapeless carriage and an alphameric print speed of 2000 lines per minute is available. The tapeless carriage decreases operator intervention by eliminating carriage tape loading and unloading.

The new 3803/3420 Magnetic Tape Subsystem is attachable. Models 3, 5, and 7 of the 3420 Magnetic Tape Unit, with data rates of 120 KB, 200 KB, and 320 KB, respectively, at 1600-BPI recording density, are provided. Phase-encoded recording, which automatically corrects all single-bit read errors in-flight, is used. This new tape subsystem offers improved price performance; Dual Density and Seven-Track features for compatibility with, and conversion of, 2400-series tape volumes; greatly reduced operator handling through implementation of such features as automatic tape threading and cartridge loading; lower cost tape switching than is currently provided; and enhanced reliabilty, availability, and serviceability features.

- Extensive hardware and programming systems error recovery and repair features are provided to improve system reliability, availability, and serviceability.

- Compact physical design reduces Model 165 CPU and processor storage space requirements. The Model 165 CPU has three times the number of circuits as a Model 65, in excess of 75,000 versus 25,000, yet a 512K or 1024K Model 165 requires the same amount of space as a 512K Model 65.

As the highlights indicate, Model 65 and 75 users have a broader range of Model 165 configurations to choose from than before when tailoring a growth system with improved throughput and expanded capabilities.

Specifically, the Model 165 offers the following advantages over Models 65 and 75:

Larger Processor (Main) Storage Sizes

Storage sizes of 512K, 1024K, 1536K, 2048K, and 3072K are provided. The Model 65 offers a maximum of 1024K. Larger Model 165 storage sizes are available at smaller cost increments, and additional storage can contribute significantly to system performance and capabilities.

The addition of more processor storage provides the Model 165 user with the ability to:

- Execute more jobs concurrently, including new application and integrated 7000-series emulator jobs
- Add and expand applications, such as graphics, teleprocessing, time sharing, remote job entry, and data based, that require large amounts of storage
- Use higher level language translators and linkage editors that provide more functions and execute faster
- Execute larger processing programs without the necessity of overlay structures
- Allocate more processor storage to language translators and sorts to improve their execution speed
- Use more and larger I/O buffers to speed up input/output operations and optimize use of direct access storage space
- Include system generation options that improve control program performance and support additional functions

Expanded Channel Capabilities

The fast internal performance of the Model 165, together with expanded use of multiprogramming, requires that more data be available faster.

The Model 165 offers more and faster channels than the Model 65, twelve instead of seven, and 1.5 MB and 3.0 MB data rates on the 2880 in addition to a 1.3 MB rate on the 2860.

The channel features of the Model 165 provide:

- A significantly higher attainable aggregate data rate than provided by the Model 65 to balance the high performance capabilities of the Model 165 CPU. A twelve-channel Model 165 system can support I/O configurations with an aggregate data rate in excess of 9 MB.
- Attachment of high-speed direct access devices, such as 3330 and 2305 facilities, that will increase I/O throughput
- Channel throughput increase capabilities via use of block multiplexing with rotational position sensing to improve effective data transfer rates

Faster I/O Devices with Increased Data Capacity

The 3330 and the 2305 facilities offer significantly faster data access than the 2314 facility and 2301 Drum Storage because of higher data transfer rates, faster rotation, and new features. Rotational position sensing and multiple requesting used with block multiplexing

can improve I/O throughput by making more efficient use of channel time. These direct access facilities also offer higher availability through use of advanced hardware-only and program-assisted error correction features.

The 3330 facility provides high capacity and fast access for less cost per bit. It is a growth device for the 2314 facility and the 2321 Data Cell Drive that offers increased price performance. The 3330 facility is designed to be used in every area in which direct access storage is needed, for example:

- As a system residence device and for program library storage
- In teleprocessing applications for message queuing and residence of online applications data
- In online, data-based applications, such as management information systems, airline reservations, etc.
- In time-sharing (or interactive) environments as a swap device and for online work storage (for program and data residence)
- As high-speed work storage for sorting, assembling, and link editing
- For residence of data indexes, such as for ISAM data sets

The 2305 facilities offer larger capacity and faster access than the 2301 drum. For Model 165 users, the 2305 facilities can contribute significantly to system throughput improvements when used:

- As system residence devices
- In time-sharing environments, as a swap device and for program and data residence
- As high-speed work storage and for residence of data indexes

SUMMARY

The combination of new and improved hardware and input/output facilities, enhanced operating systems support, integrated emulation, and increased system availability provided by the Model 165 offers Model 65 and 75 users expanded computing capabilities without the necessity of a large conversion effort. Little or no time need be spent modifying operational System/360 code or 7000-series programs currently being emulated. Existing CPU-bound programs can execute faster because of the increased internal performance of the Model 165, while I/O-bound programs can benefit from the use of more storage, more channels, faster I/O devices, and block multiplexing.

The increased power and new functions of the Model 165 provide the base for expanded applications growth and penetration of previously marginal application areas. The increased price performance of the Model 165 offers the user the opportunity to widen his data processing base for comparatively less cost and the Model 165 can be an integral part of a growth plan to the higher performance Model 195.

SECTION 10: ARCHITECTURE AND SYSTEM COMPONENTS

10:05 ARCHITECTURE DESIGN

The basic design objectives embodied in System/370 Model 165 architecture provide System/360 Model 65 and 75 users and 7000-series emulator users with a growth system that incorporates improvements and additions to System/360 architecture. The Model 165 provides new system capabilities, performance improvements, and features to increase system reliability, availability, and serviceability. This has been achieved under the following conditions:

- System/370 Model 165 architecture is upward compatible with that of System/360 models so that most user programs written for System/360 will run efficiently on the Model 165 without modification.
- Programming systems support for the Model 165 is based on that provided for System/360 models, namely on OS MFT and MVT.
- Most currently announced System/360 I/O devices will operate on the Model 165. (See Section 20:05 for a list of the I/O devices that cannot be included in a Model 165 configuration.)
- The open-ended design characteristic of System/360 has been preserved and extended in System/370.

As a result of the architecture design criteria used for this new system, all programs written for System/360 (Models 25 and up) will operate on a System/370 Model 165 with a comparable hardware configuration, with the following exceptions:

1. Time-dependent programs
2. Programs using machine-dependent data such as that which is logged in the machine-dependent logout area. (OS SER and MCH and DOS MCRR error-logging routines for System/360 models will not execute correctly.)
3. Programs that use the ASCII mode bit in the PSW
4. Programs that depend on the nonusable lower processor storage area being smaller than 1512 bytes. This area can be reduced to 512 bytes by moving the CPU logout area. (See Section 50.)
5. Programs deliberately written to cause certain program checks
6. Programs that depend on devices or architecture not implemented in the Model 165, for example, the native file of the Model 44, relocation implemented in the Model 67, etc.
7. Programs that use model-dependent operations of the System/370 Model 165 that are not necessarily compatible with the same operations on System/360 models

Note that these are the same types of restrictions that exist for compatibility among System/360 models.

The major elements of the Model 165 computing system are illustrated in Figure 10.05.1. Each component and its new features are discussed

in the subsections that follow. Programming systems support of these features is covered in Section 30. Reliability, availability, and serviceability (RAS) hardware features are mentioned only briefly. A full discussion of both hardware and programming systems RAS facilities is contained in Section 50.

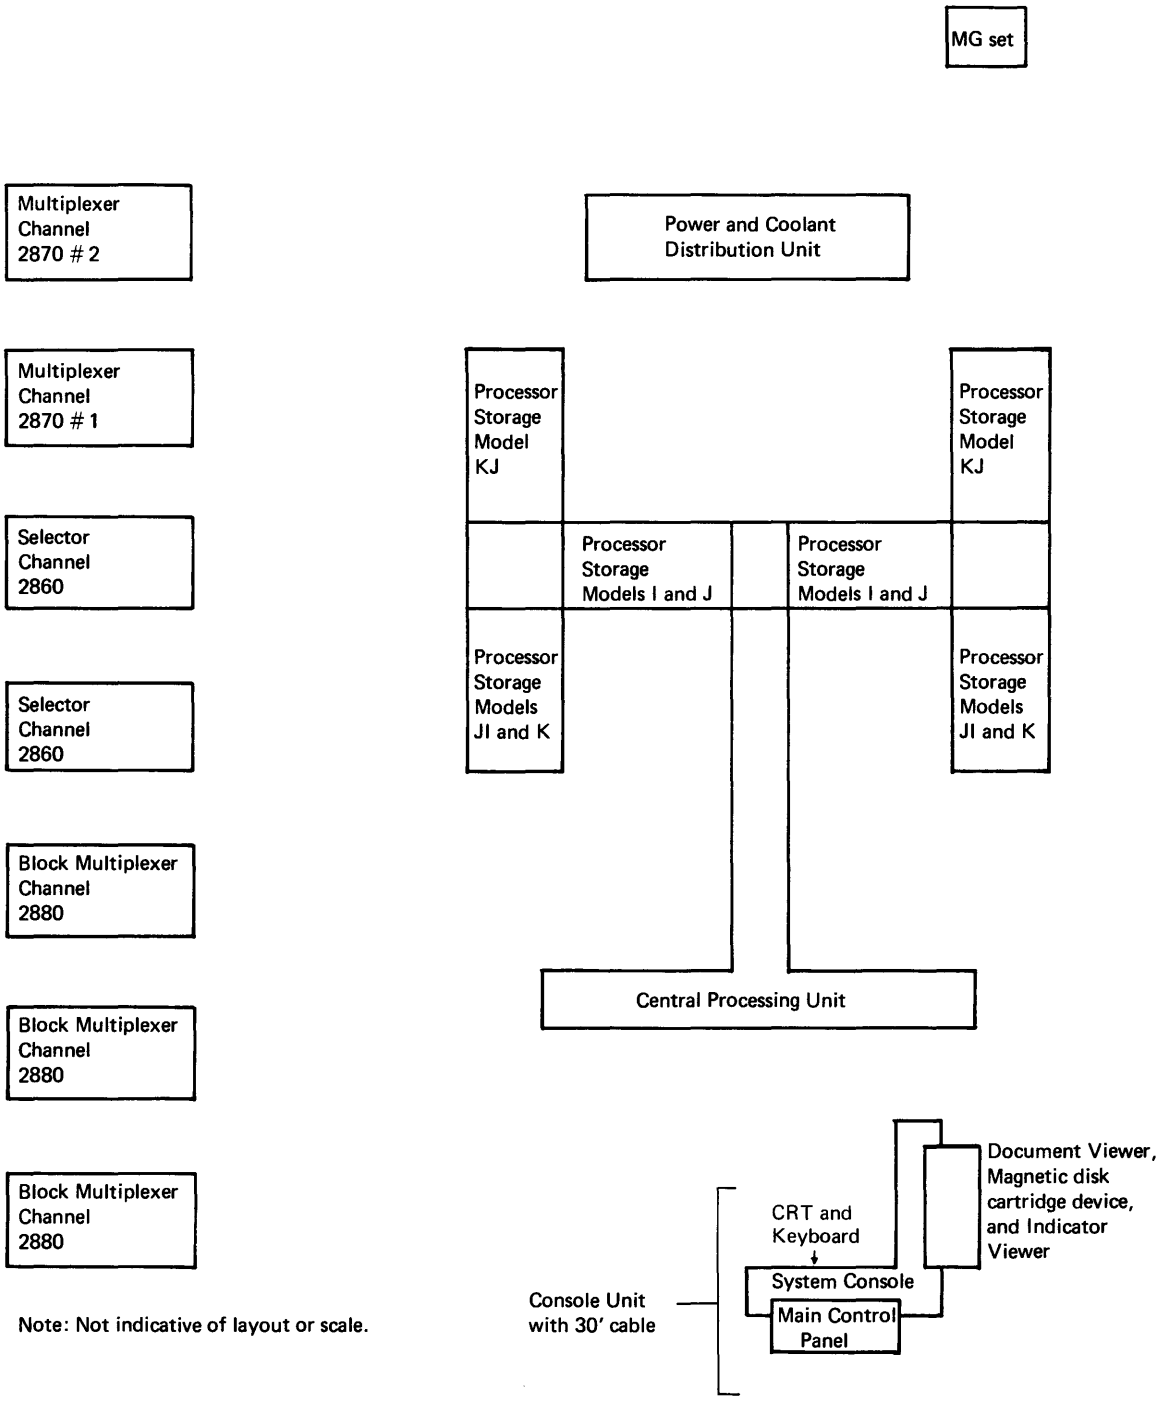

Figure 10.05.1. System/370 Model 165 system elements

10:10 THE CENTRAL PROCESSING UNIT (CPU) AND THE SYSTEM CONSOLE

CENTRAL PROCESSING UNIT

The CPU contains all the elements necessary to decode and execute the instructions in the System/370 Model 165 instruction set and, optionally, those in the hardware compatibility feature required by one of the three 7000-series emulator programs.

The CPU has an 80-nanosecond cycle time and an 8-byte-wide data path. Extensive parity checking is performed in the CPU to insure the validity of the data being used. All data transfer, logical, and arithmetic operations are checked. Automatic hardware retry of most failing CPU operations, without programming assistance, is provided as an availability feature and is discussed in the RAS section.

Among the major elements in the CPU are the instruction unit, the execution unit, local storage, and control storage.

Instruction and Execution Units

The faster internal performance of the Model 165 is due in part to the use of more concurrence in CPU operations than is implemented in the Model 65. The Model 165 CPU contains an instruction unit and an execution unit that overlap instruction fetching and preparation with instruction execution. The Model 165 instruction unit is controlled by logic circuits and can process several instructions concurrently while the execution unit is executing a single instruction.

The instruction unit prefetches instructions (maintaining them in sequence), decodes instructions, calculates addresses, prefetches instruction operands, and makes estimates of the success of conditional branches. When a conditional branch is encountered, the instructions immediately following the branch and those located at the branch address are prefetched and placed in separate instruction buffers within the instruction unit. Two 16-byte instruction buffers are used. This insures the availability of prefetched instructions whether the branch is taken or not.

The execution unit is microprogram controlled and can execute one instruction at a time. It has the capability of processing a new instruction every cycle. Emphasis is placed on optimizing fixed binary and floating-point arithmetic operations. A 64-bit parallel adder is used to perform binary and floating-point arithmetic, while an 8-bit serial adder is used in the execution of packed decimal arithmetic.

An imprecise interrupt occurs on a Model 165 only if an attempt is made to store data at an invalid storage address or at a storage-protected location.

Local Storage and Control Storage

Local storage contains the general purpose and floating-point registers and has a read or write cycle time of 80 nanoseconds. It can be accessed by four sources and written into from one source simultaneously.

Model 165 control storage consists of a capacitor read-only storage (ROS) and a monolithic writable control storage (WCS), both of which have an 80-nanosecond cycle time. ROS and WCS contain all required microcode for a specific system configuration. ROS contains the microcode necessary to execute the majority of the Model 165 instructions and some specialized routines. WCS contains the

microprogramming required to handle the balance of the instruction set for the Model 165 and other optional features, such as a 7000-series compatibility feature.

WCS is also used to house diagnostic routines. The use of some writable control storage in the Model 165 in addition to ROS allows nonresident diagnostics to overlay each other. Thus more extensive diagnostics can be provided without the necessity of adding more control storage.

During a power-on sequence, WCS is automatically loaded with system microcode from a removable magnetic disk cartridge contained on a device in the console unit. The magnetic disk cartridges sent to an installation will be tailored to include the microcode required by the optional features included in the system configuration.

Program States and System Interrupts

The program states in which the Model 165 is operating are reflected in the current program status word (PSW) and in new CPU status indicators, called control registers, located in the CPU. Up to 16 control registers, 0-15, can be addressed; however, only 4 are implemented in the Model 165. They are program addressable when the CPU is in the supervisor state. A control register can be set with the new LOAD CONTROL instruction, and its contents can be placed in processor storage with the STORE CONTROL instruction. Additional status indicators contained in control registers are required in order to support new system functions.

The contents, layout, and function of fixed locations 0-127 in System/370 models are identical to these locations in System/360 models with one exception. Bit 12 in the PSW, which sets EBCDIC or ASCII mode in System/360 models, is not used for this purpose in the Model 165. (It must be set to zero.) ASCII mode is not implemented in the Model 165, nor was the mode bit supported by IBM programming systems for System/360 models, since the expectation that System/360 ASCII-8 would become the ASCII standard has not been borne out.

The implementation of the machine check level of interruption in the Model 165 has been altered considerably from its implementation in Models 65 and 75 in order to enhance system availability (see Section 50). However, the other four interrupt levels operate in the same manner on the Model 165 as on Models 65 and 75.

CPU Features

Significant features of the Model 165 CPU are the following:

Expanded Instruction Set

The standard instruction set for the System/370 Model 165 is a superset of that provided for System/360 Models 65 and 75. It consists of the System/360 instruction set plus new instructions that support System/370 architecture and provide additional functions. The Model 165 standard instruction set includes all general purpose and I/O instructions and all binary, decimal, floating-point, and extended precision floating-point arithmetic instructions. Storage protect and time of day clock instructions are also standard. The new STORE CPU ID instruction permits a program to determine the model upon which it is operating and provides the system serial number. The new STORE CHANNEL ID instruction can be used to identify the types of channels present in the system. Other new standard instructions are:

- Extended Precision Floating Point

  Precision of up to 28 hexadecimal digits, approximately equal to 34 decimal digits, is provided by the extended precision data format. Extended precision is achieved by using two doublewords (16 bytes) to represent an extended precision floating-point number instead of using one doubleword as is done in long form representation. Fourteen hexadecimal digits, or up to 17 decimal digits, of precision is provided by the long floating-point format.

  Seven floating-point instructions are included in the extended precision feature. They provide addition, subtraction, and multiplication operations for extended precision data, using a pair of floating-point registers, and the ability to round from long to short form or from extended to long form.

- General Purpose Instructions

  Six general purpose instructions, which will be of benefit to both control and processing program performance, have been added to the Model 165 standard instruction set.

  SHIFT AND ROUND DECIMAL, using a single instruction, provides right or left shifting of packed decimal data. This instruction can save 6 to 18 bytes of instruction storage and instruction execution time for each decimal shift and round operation performed in commercial processing.

  MOVE LONG provides for the movement of up to 16 million bytes from one location in storage to another with a single instruction, thereby removing the current limitation of 256 bytes per move. This instruction can eliminate the necessity of multiple move instructions or the inclusion of move subroutines. The format and operation of MOVE LONG facilitates efficient record blocking and deblocking, field padding, and storage clearing, operations frequently performed in commercial processing.

  COMPARE LOGICAL LONG can be used to compare logically two fields of up to 16 million bytes in length, thus removing the current 256-byte limit on byte compares. In addition, when an unequal compare occurs, the two characters that caused the inequality are identified.

  The MOVE LONG and COMPARE LOGICAL LONG instructions are interruptable. Thus, when an I/O operation terminates during their execution, the interrupt can be taken, and the channel is not held up awaiting termination of what might be a lengthy move or compare.

  COMPARE LOGICAL, STORE, and INSERT CHARACTERS UNDER MASK instructions provide byte addressability within the general purpose registers and permit nonword-size data that is not on a word boundary to be compared to data in a register, loaded into a register, and stored from a register. These three instructions can be of most benefit to control program programmers, to compiler writers, and to others who must manipulate processor storage addresses.

High-Speed Multiply Feature (Optional)

This feature substantially increases the internal performance of fixed- and floating-point multiply operations. With use of this feature a long precision floating-point multiply operation improves from 1.81 to .61 microseconds. A fixed-point multiply operation improves from .78 to .42 microseconds.

Architecture Implementation Alterations

Two alterations have been made to the system action taken on a Model 165 during the execution of certain instructions common to both System/370 and System/360 models. The first involves all instructions that check the validity of operands involved in packed decimal operations. On the Model 165, an invalid sign in an operand causes the instruction to be suppressed (never executed) rather than terminated during execution as is done on System/360 models. Suppression, rather than termination, of an instruction when an invalid sign occurs insures that the data fields involved remain unchanged. Therefore, when a program check occurs, a routine can be executed that inspects the field that has the invalid sign.

For example, when an invalid sign results from packing an entirely blank field, the sign can be corrected by programming, and transaction deletion or program termination is avoided.

The second alteration concerns the recognition of a storage protection exception during the execution of an EDIT or an EDIT AND MARK instruction. On a Model 165 a protection exception always occurs when a pattern character is fetched from a location protected for storing but remains unchanged during the edit operation. This change eliminates unpredictable system operation during editing on a Model 165. The occurrence of a protection exception for the situation described is model-dependent for System/360 models.

Interval Timer (Standard)

The interval timer in decimal location 80 in fixed processor storage of a Model 165 has a resolution of 3.33 ms instead of the 16.6 ms resolution (with 60-cycle power) implemented for the standard timer on the Model 65. Its maximum time period remains 15.5 hours. The higher resolution of this interval timer will eliminate many of the problems encountered in accounting routine accuracy caused by task execution durations that are less than the 16.6 ms resolution.

Time of Day Clock (Standard)

This new clock is a binary counter of 52 bits with a cycle of approximately 142 years. It is updated every microsecond. Two new instructions (SET CLOCK and STORE CLOCK) are provided to set the time and to request that the current time be stored in a specified doubleword of processor storage. The time can be set only when the CPU is in supervisor state and only when the clock security switch on the system console panel is in the enable set position.

The time of day clock can be used for more accurate time stamping than the interval timer. Accurate time of day can be maintained because during normal operations the clock stops only when CPU power is turned off. The interval timer cannot be as accurate as the clock for time of day maintenance because it is not updated when the system is in the stopped state, and its updating may be omitted under certain conditions of excessive system activity. The 15.5-hour cycle time of the interval timer is also a restriction. The time of day clock better answers the timing needs of teleprocessing and real-time applications.

CPU Cooling

The heat generated by the logic boards in the Model 165 CPU and its associated power frames is removed by forced air and a closed-loop water circulation system. Use of a liquid coolant in addition to air is required because of the amount of heat generated by the densely packed circuits in the CPU.

The user must supply 30 gallons of cooled water per minute (45° to 60° F.) to the coolant distribution unit (CDU), which is housed in a stand-alone frame that also contains power and the power distribution unit. Water is supplied to the CDU in pipes under the raised floor. The chilled water entering the CDU is used to control the temperature of the internal water that passes through the CPU. That is, the user-supplied water does not enter the closed-loop system of the CPU. (See Figure 10.10.1.) The CDU houses the necessary controls to maintain the proper temperature in the closed-loop system. The user must supply the controls to maintain the temperature of the chilled water supplied to the CDU.

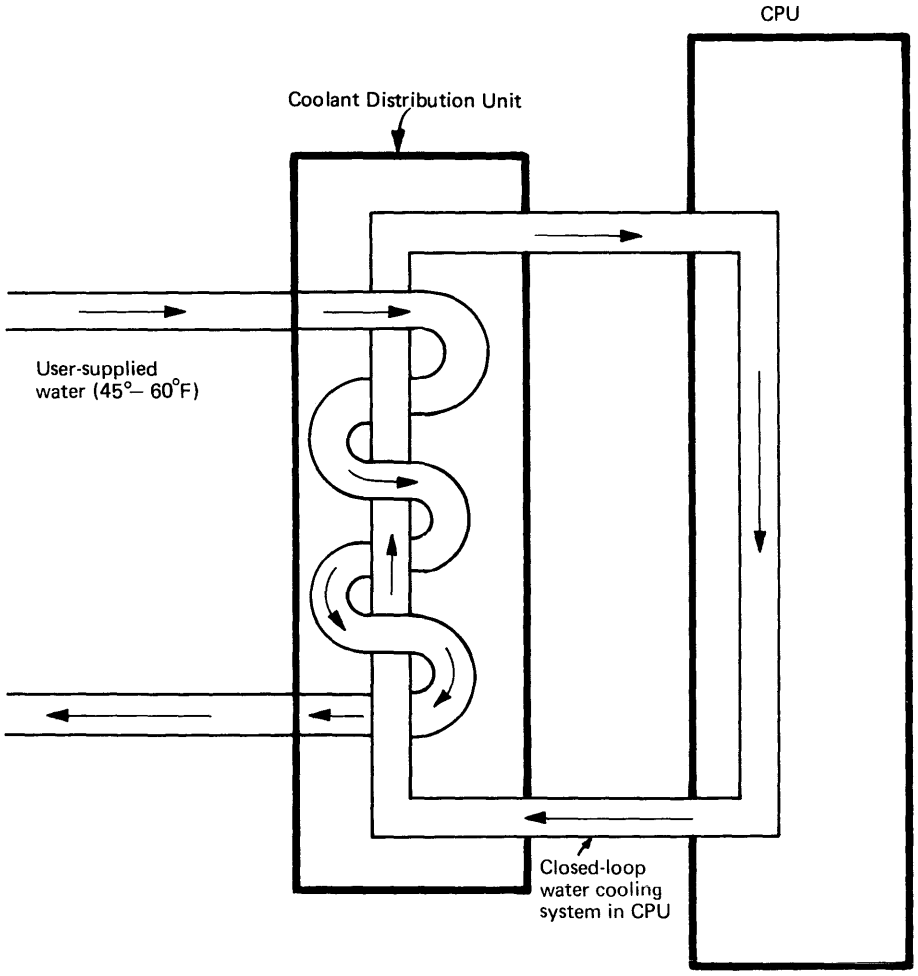

Figure 10.10.1. Conceptual flow of the water cooling system in the Model 165

The use of water as the cooling liquid offers several advantages. First, it is readily available. The chilled water normally supplied for air conditioning is acceptable. Second, water offers safety features. Low pressure is required and the cooling system can operate at room temperature, thereby eliminating problems with condensation. Last, a simplified circulation system suffices, with relatively few moving parts and less exposure to leaks. Pipe components within the CPU are separated easily, and the valved connections close off automatically to prevent water from escaping.

Physical planning for Model 165 installation should insure that arrangements are made to provide the required water. (See System/370 Installation Information, Physical Planning, GA22-6971 for more details.)

Motor Generator Set

The motor generator (MG) set is the converter unit that provides the power required by the Model 165 CPU. It takes 60 Hz (cycle) power from the building electrical distribution system, converts it to 415 Hz power, and supplies it to the CPU.

Unlike the MG sets provided for the 7000-series, which were small enough to be housed within the covers of the system itself, the MG set (including the required starter) for the Model 165 is a stand-alone unit approximately six feet long, weighing 3000 lbs. (See System/370 Installation Information, Physical Planning for more details concerning MG set size and installation requirements.)

The MG set should be ordered at the same time as the Model 165, with delivery up to two months prior to system installation. While IBM does not manufacture MG sets, a procedure is established for ordering the required MG set through IBM.

3066 SYSTEM CONSOLE

A stand-alone system console unit, which can be located up to approximately 25 feet away from the CPU, is required for the Model 165. This unit contains the system control panel (buttons, switches, lights, etc., required for system operation), as well as standard advanced console features: a cathode ray tube and keyboard, a microfiche indicator viewer, a microfiche document viewer, a processor storage configuration plugboard, a system activity monitor, and a device for loading microcode and diagnostics. Certain of these features are included to enable the customer engineer to detect CPU malfunctions more rapidly than would otherwise be possible.

Installation of the optional Remote Operator Console Panel Attachment feature on a Model 165 permits duplication of the operator control panel section of the system console unit on a 2150 Console or a 2250 Display Unit Model 1.

The cathode ray tube (CRT) unit provided contains a 4K buffer. In normal mode, the CRT and alphameric keyboard are designed to be used as the operating system operator's console to provide the rapid message display required by large systems like the Model 165. Two operator warning indicators are provided. A switch setting determines whether a visual indicator or an audible alarm with volume control is used. In CE mode, the CRT and a set of 16 data keys can be used to display and manually alter processor storage, the general purpose and floating-point registers, the channel buffers, or the CPU buffer address array. Data is entered in hexadecimal by means of the data keys, each of which represents a hexadecimal digit.

The microfiche indicator viewer is included in the console unit for display of status and control indications. The display utilizes a "framing" concept: the customer engineer uses a "frame selection switch" to select any one of ten images. Each image consists of up to 240 indicators, the current value of which is then displayed in incandescent lights along with the appropriate microfiche image that labels the lights. In this way up to 2400 status and control triggers may be displayed, yet the console requires only 240 incandescent lights.

A microfiche document viewer is included in the console unit for display of customer engineer reference materials. It is expected that more than half of all Model 165 CE reference materials will be supplied in microfiche form.

The processor storage configuration plugboard permits automatic generation of the required processor storage configuration (units available to the operating system, their addressing, interleaving, etc.) when power is brought up or upon depression of an "enter configuration" push-button. (See Section 10:15 for further information.) The exact processor storage configuration in effect at any given time can be verified by displaying the appropriate image on the indicator viewer.

The system activity monitor is capable of displaying the average activity of the CPU and the channels via the system activity meter located in the main control panel in the system console. The activity to be displayed is selected via a set of switches. The following can be displayed, one measurement at a time.

1. Percentage of compute time relative to the:

   a. Total time CPU is active
   b. Time CPU is active in supervisor state
   c. Time CPU is active in problem state

2. Percentage of time the CPU and one or more of the selected channels are concurrently active (to determine CPU and channel activity overlap)

3. Percentage of time when all of the selected channels are busy concurrently (to determine channel activity overlap)

4. Percentage of time when a single selected channel or one or more of the selected channels are busy (to determine channel activity)

5. Percentage of time any of the selected channels and the CPU are busy accessing the selected storage protect key area of processor storage

A device that provides read-only storage for WCS microcode and system diagnostics is contained in the system console unit between the indicator viewer and the document viewer. This device reads interchangeable magnetic disk cartridges. As already indicated, this device is used for WCS loading. It is also a debugging tool for the customer engineer who will load diagnostics from it.

10:15 STORAGE

The Model 165 has a two-level memory system - a very high-speed buffer storage backed by a large processor (main) storage. Prior to this time, such a concept was implemented only in very large-scale, high-internal-performance systems, such as System/360 Models 85 and 195. The Model 165 has four-way doubleword interleaved processor storage with a two-microsecond cycle time, buffered by 8K or 16K of high-speed monolithic buffer storage with an 80-nanosecond cycle time.

The use of interleaving and a two-level memory system drastically reduces the effective storage cycle of the Model 165 and greatly contributes to the fact that the internal performance of the Model 165 is approximately two to five times that of the Model 65.

PROCESSOR (MAIN) STORAGE

Processor storage is available in the following increments:

Model	Capacity
I	512K
J	1024K
JI	1536K
K	2048K
KJ	3072K

In order to achieve interleaving in the Model 165, the processor storage present in the system is divided into logical memories within storage boxes. A logical memory is defined as that portion of storage that can operate independently from all other storage. A logical memory in the Model 165 is 128K or 256K and there are two per storage box.

Depending on the storage size, the total number of logical memories present is 4 in two boxes for Models I and J, 8 in four boxes for Models K and JI, or 12 in six boxes for Model KJ, as shown in Figure 10.15.1. (Model JI, not shown, consists of four 128K and four 256K logical memories in two boxes.) However, only four logical memories can be active at one time, thus achieving four-way interleaving. As illustrated in Figure 10.15.1, for a 512K system, four consecutively addressed doublewords are spread across logical memories so that 32 consecutive bytes can be accessed concurrently (one doubleword from each of four logical memories, requested one 80-nanosecond cycle apart). That is, processor storage location 0 is in logical memory 0, location 8 is in logical memory 1, location 16 is in logical memory 2, and location 24 is in logical memory 3. Processor storage location 32 is in logical memory 0, and the address distribution sequence continues through all available storage locations.

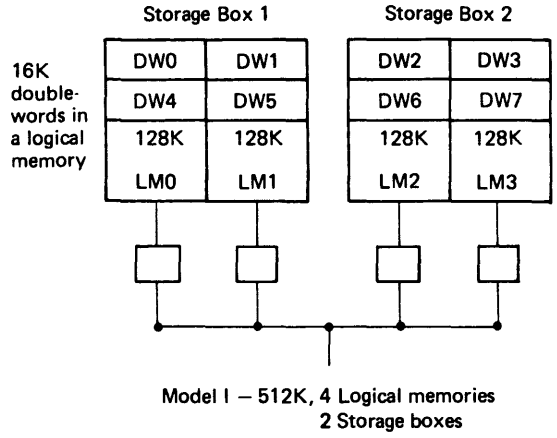

Model I – 512K, 4 Logical memories
2 Storage boxes

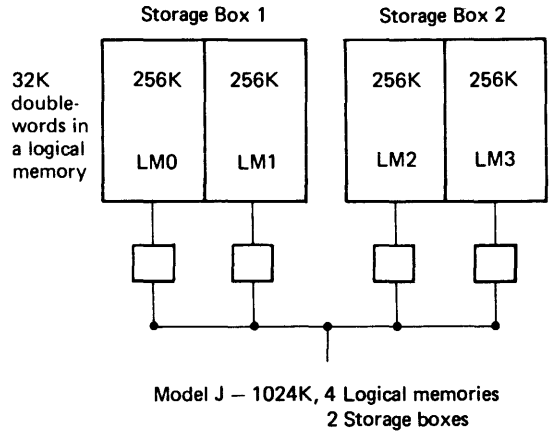

Model J – 1024K, 4 Logical memories
2 Storage boxes

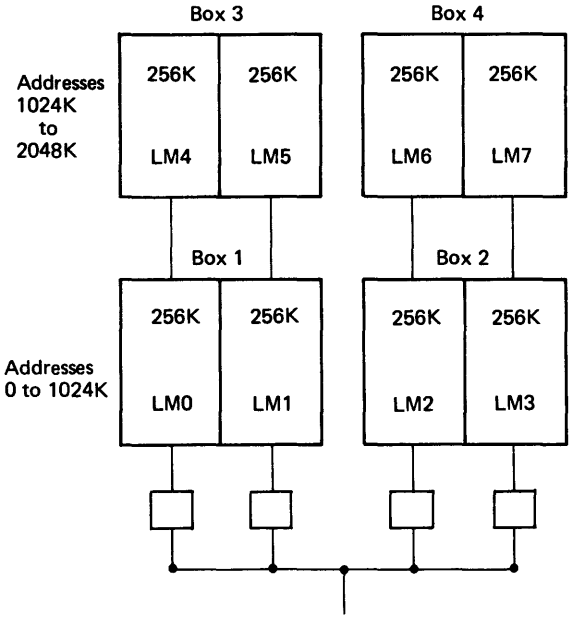

Model K – 2048K, 8 Logical memories
4 Storage boxes

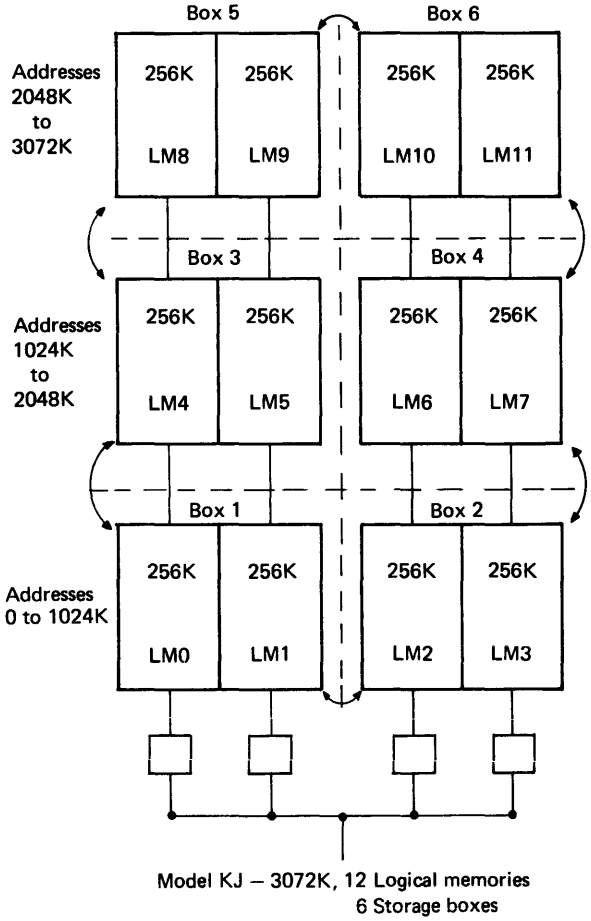

Model KJ – 3072K, 12 Logical memories
6 Storage boxes

Figure 10.15.1. Model 165 processor storage configurations

Processor storage can be accessed concurrently by any combination of one or more channels and the CPU for a total of four unique logical memory requests. Contention arises when two or more components attempt to access the same logical memory simultaneously. In addition, requests made to a busy logical memory are delayed until the memory becomes free. The storage control unit (SCU), illustrated in Figure 10.15.2, controls all storage references by components and schedules requests according to a priority scheme. The channels have priority over the CPU for simultaneous requests and there is a definable priority among channel positions. Thus, the SCU ensures that an effective storage access of 32 bytes in two microseconds is achieved for the system whenever possible.

Storage access times stated in this guide are obtainable, assuming the four logical memories required are free when requested and accessed concurrently with no interference from other components.

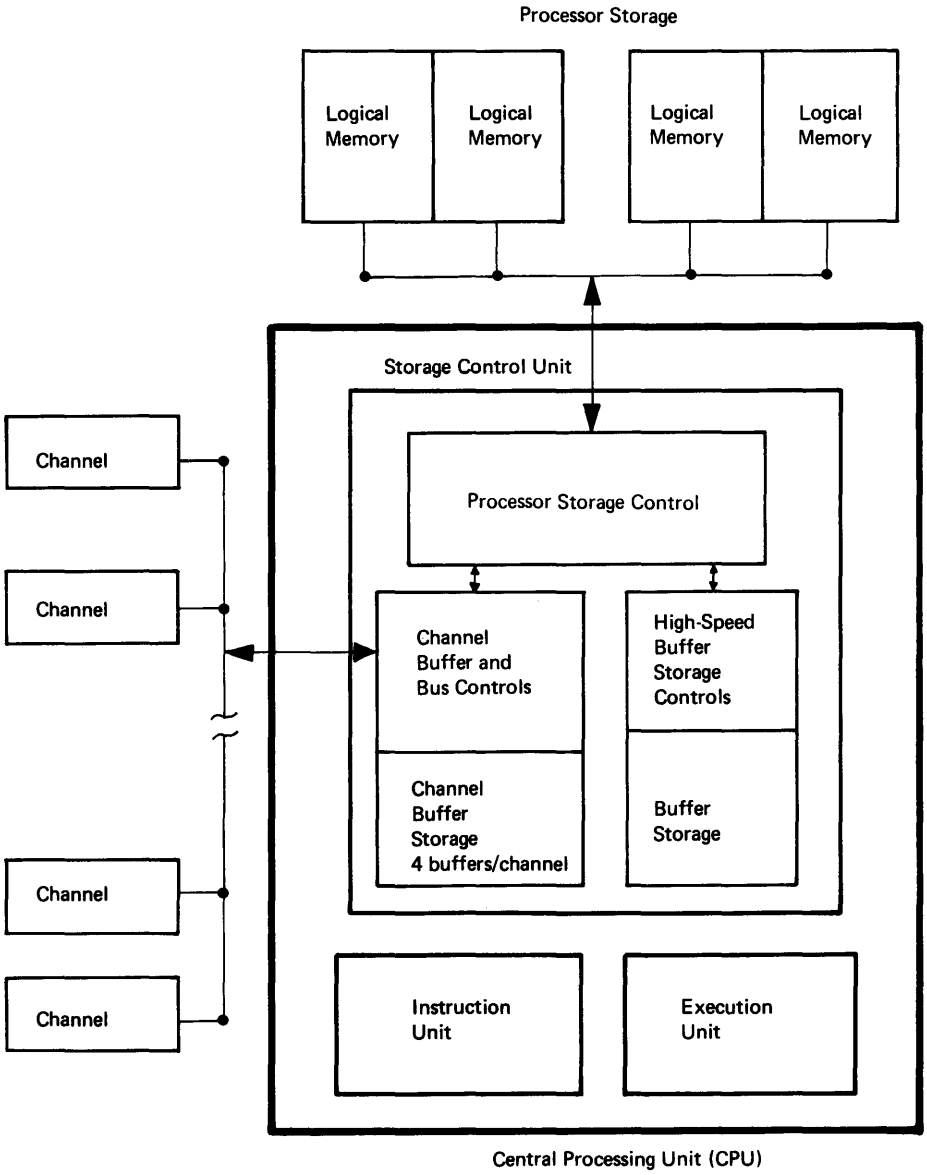

Figure 10.15.2. Model 165 storage components and controls

Error checking and correction (ECC) hardware provides automatic detection and correction of all single-bit processor storage errors and detection, but not correction, of all double-bit and most multiple-bit errors. The ECC feature is discussed fully in the RAS section.

The Model 165 also supports a byte boundary alignment facility for processor storage. The presence of the byte-oriented operand function allows the storage operands of unprivileged instructions (RX and RS formats) to appear on any byte boundary without causing a specification program interrupt. Without this facility, operands must be aligned on integral boundaries, that is, on storage addresses that are integral multiples of operand lengths. Byte orientation is standard and does not apply to alignment of instructions or channel command words (CCW's).

Use of byte alignment in a program degrades instruction execution performance. However, byte orientation can be used effectively in commercial processing to eliminate the padding bytes added within records and to blocked records to insure binary and floating-point field alignment. The smaller physical record that results from the elimination of padding bytes requires less external storage and increases effective I/O data rates. I/O-bound commercial programs, in which throughput is in almost direct proportion to the I/O data rate, can achieve performance improvement by using byte alignment for binary and floating-point data.

A program written to use byte boundary alignment will not necessarily run on a System/360 model that does not have the feature. Therefore, programs that are to run on both the Model 165 and on System/360 models without byte orientation should be written to adhere to integral boundary rules.

Processor Storage Reconfiguration

If a processor storage box develops a malfunction, it can be configured out of the system by use of the storage configuration plugboard in the system console. Then the operating system can be re-IPLed and the system can continue operating with reduced available storage. The configuration indicated by the plugboard is established during a power-on sequence or a system reset operation.

The user has the ability to remove one or more storage boxes from the operative system and reconfigure the addressing of the remaining boxes to achieve consecutive storage addressing. Interleaving is reduced from four-way to two-way if the configuration consists of an odd number of boxes. Therefore, four-way interleaving can be maintained in systems with four or six processor storage boxes by removing a pair of boxes instead of the malfunctioning box only. A pair must be removed in a 3072K (six-box) system as a 2560K configuration is not supported. Serial operation is possible also and will be used primarily by customer engineers.

The configuration panel is relatively simple to use. The operator inserts plugs into the appropriate holes in the panel to describe the processor storage configuration: number of boxes (one to six) using up to three plugs, box addressing sequence (box reversals) using up to two plugs, and interleaving (four-way, two-way, or serial) using up to two plugs.

With the few reversal combinations defined, any box can be placed in the first or last box addressing position. Assume box 2 in a six-box configuration is to be configured out of the operational system. (Refer to Figure 10.15.1, in which box numbers also indicate the sequential positioning of consecutive processor storage box addressing.) Boxes 1 and 2 would be vertically reversed in position with boxes 3 and 4, respectively, and then with boxes 5 and 6, respectively. Box

2 is then in the last (or highest) addressing position and an inline ripple (discussed below) can be performed on it. If box 1 is the malfunctioning box, a horizontal reversal of boxes 3, 5, and 1 with boxes 4, 6, and 2, respectively, puts box 1 in the highest addressing position. (The ascending addressing sequence of the boxes is then 4-3, 6-5, 2-1.)

Storage Ripples

Five storage ripple functions are implemented in Model 165 hardware as maintenance aids for use by customer engineers. (A ripple is a nonprogrammed read or write of ones or zeros through every available storage address for the purpose of locating a malfunction.) A ripple is provided for ROS, WCS, local storage, and processor storage that requires dedication of the system to the ripple function. However, an inline processor storage ripple also is implemented. It can be executed on a malfunctioning storage box that has been configured out of the operational system while processing continues. (The processor storage box must be in the highest addressing position.) An inline ripple is not provided for the Model 65, which therefore requires total system unavailability during processor storage rippling. In addition, there is no reconfiguration capability for main storage boxes in uniprocessor Model 65 systems.

HIGH-SPEED BUFFER STORAGE

The increase in the internal performance of the Model 165 is achieved largely by the inclusion of a high-speed buffer storage unit. The 8K buffer is a standard feature and provides high-speed data access for CPU fetches. Installation of the optional Buffer Expansion feature permits inclusion of an additional 8K of buffer storage.

The buffer has an 80-nanosecond cycle. The CPU can obtain eight bytes from the buffer in two cycles, or 160 nanoseconds, and a request can be initiated every cycle. This compares with 1.44 microseconds (or 18 cycles) required to obtain eight bytes of data directly from processor storage. The conceptual data flow in the Model 165 is pictured in Figure 10.15.3.

Buffer storage control and use is handled entirely by hardware and is transparent to the programmer, who need not adhere to any particular program structure in order to obtain close to optimum use of the buffer. The buffer algorithm implemented in the Model 165 is very similar to that used in the System/360 Model 195. Sample job step executions have shown that in a Model 165 the data accessed by the CPU is in the buffer 95% of the time on the average.

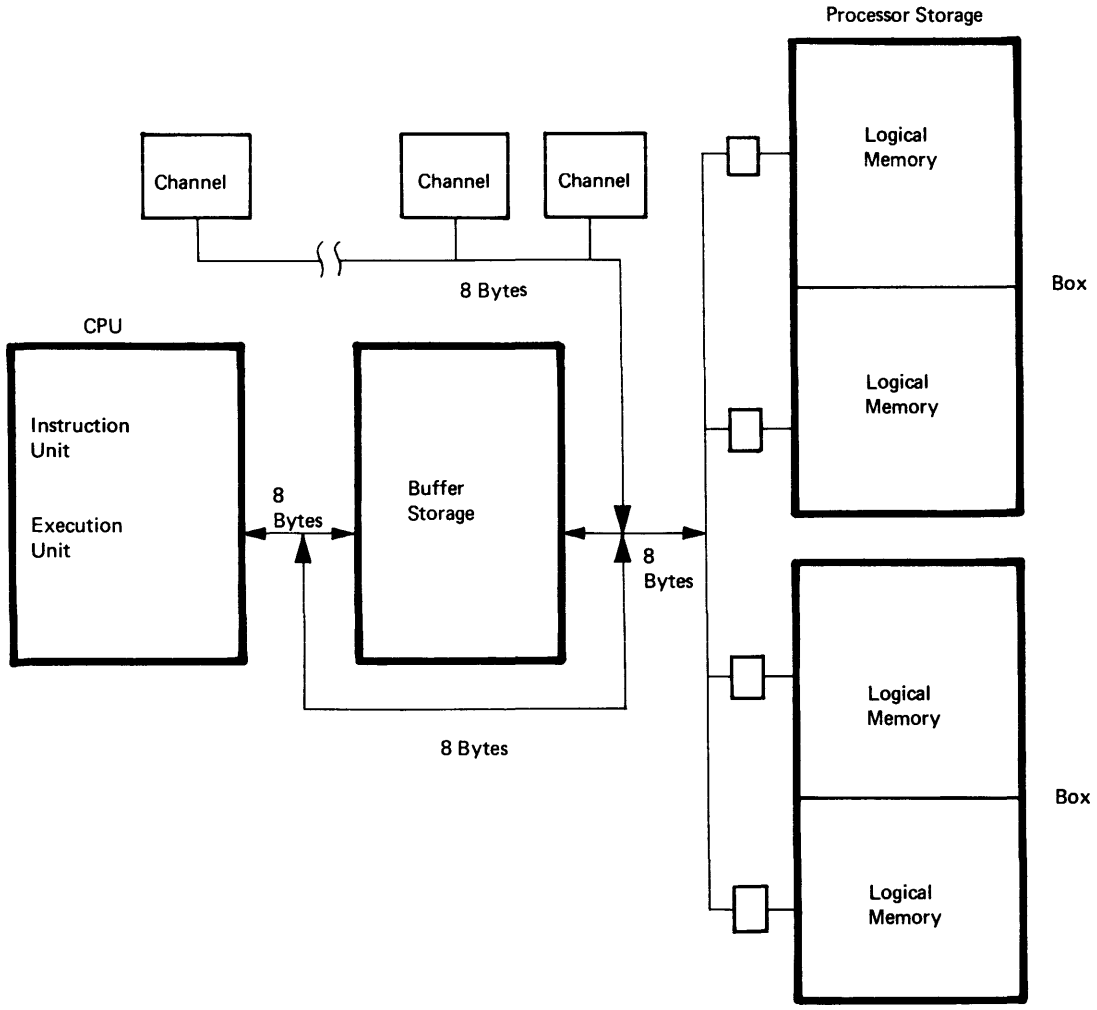

CPU fetch from buffer - 160 nanoseconds for 8 bytes (2 cycles)
CPU fetch from processor storage - 1.44 microseconds for 8 bytes (18 cycles)
Buffer fetch from processor storage - 1.44 microseconds for 32 bytes (18 cycles)
Channels to and from processor storage - 32 bytes in a 2 microsecond cycle

Figure 10.15.3. Conceptual data flow in the Model 165. Times given assume no interference.

The storage control unit (SCU) contains the high-speed buffer and controls all buffer and processor storage references made by the CPU, the channels, and manual controls. The buffer storage control portion of the SCU handles CPU to processor storage references, both fetches and stores. Parity checking is used for data verification in the buffer.

When a data fetch request is made by the CPU, buffer storage control determines whether or not the requested data is in the high-speed buffer by interrogating its address array of the buffer's contents. If the data requested is present in the buffer, it is sent directly

to the CPU without a processor storage reference. If the requested data is not currently in the buffer, a processor storage fetch is made. The data obtained is sent to the CPU. The data is also assigned a buffer location and stored in the buffer. When data is stored by the CPU, both the buffer and processor storage are updated if the processor storage location being altered is one whose contents are currently being maintained in the buffer.

The channels never access the buffer directly. They read into and write from processor storage only. When a channel stores data in processor storage, the address array is interrogated. If data from the affected processor storage address is being maintained in the buffer, the data is placed in the buffer as well as in processor storage.

The entire buffer can be disabled manually by a system console switch or via execution of a DIAGNOSE instruction. When the buffer is disabled, all CPU fetches are made directly to processor storage and effective system execution speed is reduced.

The 8K buffer is shown in Figure 10.15.4. It contains 64 columns of 128 bytes each. Every buffer column is subdivided into four blocks. A block is 32 bytes and can contain 32 consecutive bytes from processor storage that are on a 32-byte boundary. The 8K buffer can contain a maximum of 256 different blocks of processor storage data (4 blocks per column times 64 columns). A valid trigger is associated with each buffer block and is set to indicate whether or not the block contains valid data. All valid triggers are set off during system reset or IPL.

Processor storage is logically divided into the same number of columns as buffer storage: 64 or 128. While there are always four blocks in a buffer column, the number of blocks in a processor storage column varies with the size of processor storage. For example, when an 8K buffer is present, bits 21-26 of the processor storage address determine which of the 64 columns to use. As shown in Figure 10.15.4, a processor storage column consists of 512 blocks in a system with 1024K. Any of the 512 blocks in a given processor storage column can be placed in any of the four blocks in the corresponding buffer column. Figures 10.15.5 and 10.15.6 show the formats used for processor and buffer storage addressing.

The larger the processor storage size, the greater the number of storage blocks that contend for the four blocks in the same buffer column. If a 16K buffer is used instead of an 8K, the number of buffer columns is doubled from 64 to 128; the number of processor storage blocks contending for the blocks in each buffer column is thereby halved if processor storage size remains the same. (A 1024K processor storage divided into 128 columns has 256 blocks per column.)

The 16K buffer is provided for users with larger Model 165 processor storage configurations that have applications such that increased system throughput results from an increase in internal performance.

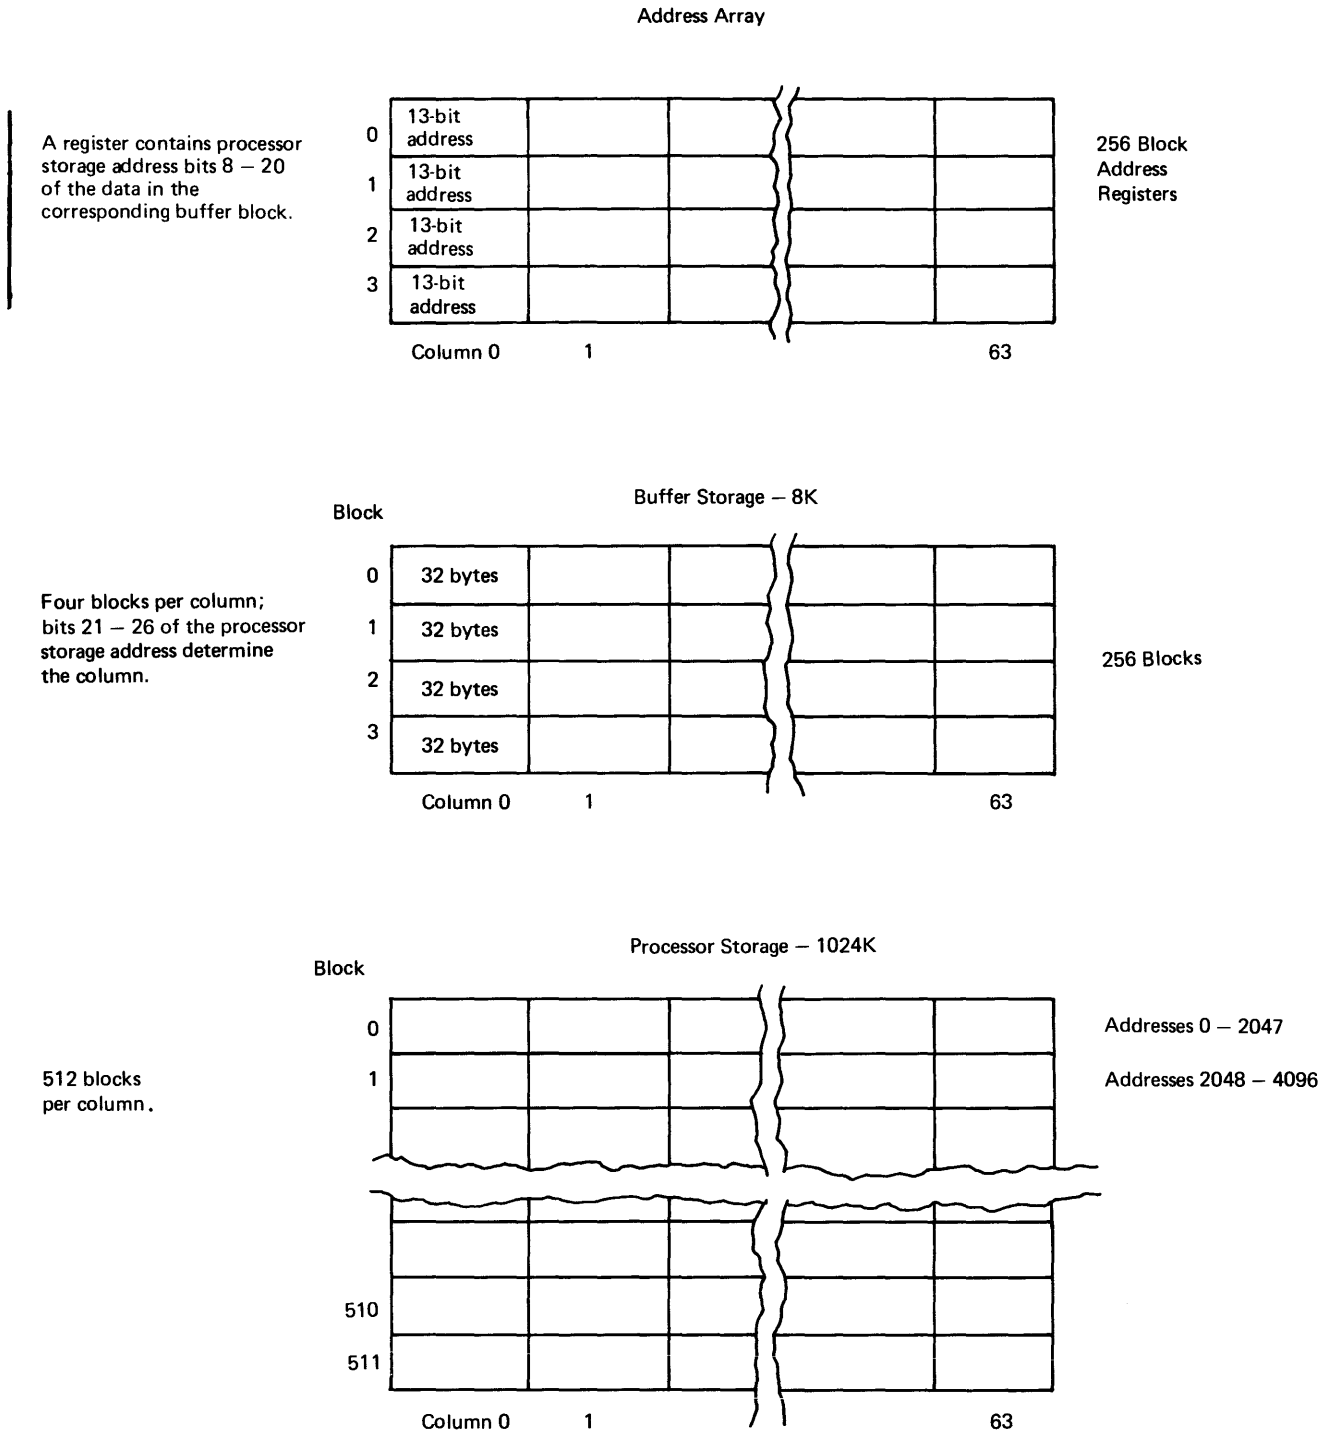

● Figure 10.15.4. 8K buffer organization

Processor Storage Address Bits

8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Bits

8-19,20	Used for address compare (16K, 8K buffer)
20-26	Used to reference 16K buffer columns
21-26	Used to reference 8K buffer columns
27-28	Used to reference doubleword within a block
29-31	Used to reference byte within doubleword

Figure 10.15.5. Processor storage address format for buffer reference

Buffer Reference Bits

0	1	2	3	4	5	6	7	8	9	10

Bits

0-1	Generated as result of address compare; this two-bit encoded field represents block in buffer column containing desired data
2-8	Used to reference column for 16K buffer (processor storage address bits 20-26)
3-8	Used to reference column for 8K buffer (processor storage address bits 21-26)
8-10	Used to reference doubleword within a block (processor storage address bits 27-28)

Figure 10.15.6. Buffer address format

Buffer contents and buffer block assignment are controlled by an address array, shown in Figure 10.15.4, and a special replacement array. The address array, like the 8K buffer, is divided into 64 columns consisting of four block address registers each such that there is a one-for-one correspondence between address array registers and blocks in the buffer. An address array block register contains the 13-bit processor storage block address from bits 8-20 of the processor storage address of the data contained in its corresponding buffer block. When a CPU to processor storage reference is made, the four appropriate address array column registers (13-bit block addresses) are interrogated to determine whether the requested data is currently in the buffer.

The replacement array is used to maintain knowledge of the activity of the data blocks within each of the buffer columns. When an 8K buffer is present, the array consists of 64 logic-controlled activity lists, one list for each column in the buffer. A list contains four entries, one for each buffer block in its column. A block's entry is placed at the top of the list for its column when the buffer block is referenced during a CPU fetch operation. This approach insures that the block used longest ago within a given column is at the bottom of the list. When a block within a buffer column has to be assigned and loaded, because the data requested by the CPU is not in the buffer, the buffer block at the bottom of a column activity list is allocated. Thus, the more active data is maintained in the buffer.

The 8K buffer operates as follows. When the CPU requests data, bits 21-26 of the data's processor storage address are used to obtain a buffer column address. The 13 high-order bits of the processor storage address are then compared to the address in each of the four block address registers in that buffer column in the address array.

If an equal compare occurs for an address in one of the registers in the address array and the valid trigger for that block is on in the buffer, the appropriate doubleword from the buffer block is sent to the CPU as determined by bits 27 and 28 of the processor storage address. A processor storage reference is not made. The referenced buffer block is put at the top of its column activity list in the replacement array.

If the desired data block is not in the buffer column interrogated, the requested data must be fetched from processor storage, sent to the CPU, and stored in the buffer. The replacement array entry for the column involved is inspected and the buffer block at the bottom of the activity list is assigned to receive the requested data from processor storage. Four processor storage references, one cycle apart, are made to obtain the 32 consecutive bytes and place them in the assigned buffer block. The valid trigger for the buffer block is set on and the 13 high-order processor storage address bits are placed in the appropriate column within the address array. The first doubleword fetched from processor storage is the one containing the data required by the CPU. It is sent to the CPU as well as to the buffer so that processing can continue as soon as possible.

Assuming no channel interference, a 32-byte buffer block can be filled from processor storage in 1.44 microseconds (or 18 cycles) with four-way interleaving utilized. One buffer block can be loaded every 2 microseconds (25 cycles), assuming no interference.

10:20 CHANNELS

GENERAL DESCRIPTION

The channels available on a System/370 Model 165 are functionally compatible with those of System/360 models. Combinations of 2870 Multiplexer, 2860 Selector, and 2880 Block Multiplexer channels can be attached to a Model 165. If the Extended Channels special feature is installed, up to twelve channels can be connected to a single Model 165 to provide an aggregate channel data rate in excess of nine megabytes, twice the rate possible on a Model 65.

The 2870 and 2860 channels that attach to the Model 165 are the same as those used on a Model 65 but have minor hardware changes (as is true when these channels are attached to System/360 Models 75 and up). When a Model 165 replaces a Model 65 or 75, the 2870/Model 165 and 2860/Model 165 features can be field installed on the existing 2870 and 2860 channels, respectively, so that the latter can be attached to the Model 165.

The standard number of addressable channels permitted on a Model 165 is seven. Any combination of one or two 2870 Multiplexer, up to six 2860 Selector, and up to six 2880 Block Multiplexer channels can be attached up to the limit of seven channels. If the optional Extended Channels feature is installed, the maximum number of each type of channel that can be included in the twelve channels permitted is:

- 2870 Multiplexer - two. (The first must have address 0, the second can be assigned an address from 1 to 6.)
- 2860 Selector - six. (Addresses 1 through 6 can be assigned.)
- 2880 Block Multiplexer - eleven. (Addresses 1 through 11 can be assigned.)

Model 165 channels are not integrated with the CPU. Channels compete with each other and the CPU only for processor storage accesses and, therefore, cause a minimum of CPU interference. A 2870, 2860, or 2880

channel contains the hardware required to control its I/O operations (channel registers, local storage, control functions, buffers, etc.). A channel interferes with the Model 165 CPU if the CPU accesses a logical memory that is busy because of a channel operation. Contention between channels and the CPU for processor storage is reduced drastically by the use of high-speed buffer storage, which eliminates CPU to processor storage fetches for approximately 95% of the fetches required.

The standard instruction set also includes a new I/O instruction called HALT DEVICE. This instruction is specifically designed to stop an I/O operation on a particular device on a multiplexer channel without interfering with other I/O operations in progress on the channel. HALT DEVICE should always be used instead of HALT I/O to stop an I/O operation on a 2880 Block Multiplexer channel.

The 2870 Multiplexer can control concurrent execution of 192 slow-speed to medium-speed devices, one with each of its 192 subchannels. Depending on the channel priority assigned, the 2870 can support a data rate of up to 110 KB. The maximum aggregate byte data rate of the 2870 Multiplexer subchannels is reduced by the inclusion of one or more selector subchannels, each of which can have up to 16 I/O devices attached. Each of the first three selector subchannels included can operate at a 180 KB rate, can handle one burst operation at a time, and reduces the aggregate byte rate of the multiplexer interface by 10 KB to 25 KB depending on the priority of the 2870 and the total I/O configuration. The fourth selector subchannel can operate at a rate of 100 KB and further reduces the maximum aggregate rate of the multiplexer subchannels by 14 KB. If two 2870 Multiplexer Channels are installed, the second 2870 can have only two selector subchannels.

There is an approved RPQ (Request for Price Quotation) for the 2870 channel that replaces all selector subchannels with multiplexer subchannels to give a 2870 a total of 256 multiplexer subchannels. This RPQ can be installed on both 2870 channels to provide a total of 512 multiplexer subchannels for a Model 165 system.

The 2860 and 2880 channels permit attachment of a wide variety of high-speed I/O devices to the Model 165. The 2860 Selector Channel handles data rates of up to 1.3 megabytes, while the 2880 Block Multiplexer can handle a 1.5-megabyte rate, which permits attachment of the 2305 Model 2 facility. The Two-Byte Interface optional feature can be installed on a 2880 to double its data transfer rate capability to 3.0 megabytes so that the 2305 Model 1 facility can be attached.

The Model 165 user has more flexibility than the Model 65 user when configuring channels. On a Model 165, channel address and channel priority are not related, as they are on the Model 65. Any channel address (0 through 11) can have any channel priority (1-12) assigned. Thus, the Model 165 user can assign channel priority by channel type in order to achieve the desired aggregate channel rate and insure that the highest speed devices are assigned the highest channel priorities.

The maximum speed of an I/O configuration possible on a Model 165 with minimal overrun exposure is a function of the I/O devices used, the channel types installed, the channel priority assignments, and the types of channel programs operating concurrently at any particular instant. Examples of configurations that will operate on the Model 165 follow.

1. A seven-channel system including one 2870 Multiplexer and six 2860 Selector Channels with a maximum aggregate channel data rate of approximately 4.25 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2860	2301 Drum	1.250
2	2	2860	2301 Drum	1.250
3	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 3 SSC3	.090
			2540, 1403, terminals on MPX	.020
				.470
4	3	2860	2314 or 3420 Model 7	.320
5	4	2860	2314 or 3420 Model 7	.320
6	5	2860	2314 or 3420 Model 7	.320
7	6	2860	2314 or 3420 Model 7	.320

2. A seven-channel system including one 2870 Multiplexer and six 2880 Block Multiplexer channels with a maximum aggregate data rate of approximately 7.8 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2880	2305 Model 1	3.0
2	2	2880	2305 Model 2	1.5
3	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 3 SSC3	.090
			2401 Model 3 SSC4	.090
			Unit record on MPX	.023
				.560
4	3	2880	3330 facility	.806
5	4	2880	3330 facility	.806
6	5	2880	3330 facility	.806
7	6	2880	2314 or 3420 Model 7	.320

3. A seven-channel system including two 2870 Multiplexer and five 2880 Block Multiplexer channels with maximum aggregate data rate of approximately 6.8 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2880	2305 Model 1	3.0
2	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 6 SSC3	.180
			2401 Model 3 SSC4	.090
			MPX	.027
				.660
3	2	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			MPX	.040
				.400
4	3	2880	3330 facility	.806
5	4	2880	3330 facility	.806
6	5	2880	3330 facility	.806
7	6	2880	2314 or 3420 Model 7	.320

4. A twelve-channel system including one 2870 Multiplexer and eleven 2880 Block Multiplexer channels with a maximum aggregate data rate of approximately 9.4 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	8	2880	2305 Model 1	3.0
2	9	2880	2305 Model 2	1.5
3	0	2870	Multiplexer aggregate	.560
4	10	2880	3330 facility	.806
5	11	2880	3330 facility	.806
6	1	2880	3330 facility	.806
7	5	2880	2314 or 3420 Model 7	.320
8	3	2880	2314 or 3420 Model 7	.320
9	4	2880	2314 or 3420 Model 7	.320
10	2	2880	2314 or 3420 Model 7	.320
11	6	2880	2314 or 3420 Model 7	.320
12	7	2880	2314 or 3420 Model 7	.320

The Model 165 with a 2 microsecond processor storage is able to achieve higher aggregate data rates than the Model 65 with a .750 microsecond main storage because Model 165 operations have been designed to optimize the use of four-way doubleword interleaving by use of more I/O buffering. Additional channel buffers are provided in the Model 165 CPU and the 2880 channel, which permit effective use of interleaving during I/O operations and help prevent data overrun when peak activity occurs.

The storage control unit contains four dedicated 8-byte buffers for each channel that are assigned when the channel is installed. In addition, each 2860 Selector Channel contains two 8-byte buffers, and a 2880 Block Multiplexer Channel contains two 16-byte buffers or four 16-byte buffers depending on its maximum rate, 1.5 MB or 3 MB.

The use of I/O buffers in the SCU allows channel requests to be made to processor storage concurrently with other channel or CPU requests when requests are for different logical memories. Additional internal buffering in the 2880 channel allows it to withstand longer wait times for requested logical memories and thus reduces channel overrun exposure.

The channel buffering scheme implemented in the Model 165 is most efficient for 2880 channels. Use is made of all four buffers available per channel, while only two buffers can be used for 2860 channels. This scheme supports overlapped processor storage requests from the same 2880 channel so that at any given channel priority position, a 2880 can sustain a higher data rate than a 2860.

A 2301 drum connected to a 2860 Selector Channel must be placed in the highest or second highest channel priority position on a Model 165. Channel address 1 or 2 must be used. For most efficient operation of 2301 drums, they should be attached to a Model 165 via a 2880 channel, since the 2301 causes much less CPU interference during data transfer operations when connected to a 2880, rather than to a 2860.

In contrast to the Model 165, the Model 65 does not contain additional buffering in its CPU and because of timing considerations cannot utilize the two-way interleaving capability of its main storage for I/O operations. Thus, in a Model 65, without the use of buffering in the SCU, the memory bus is busy for an average of 1.1 microseconds when storing one doubleword from a 2860 channel buffer. In a Model 165, with no interference, four doublewords from channel buffers in the SCU can be placed in processor storage in 2 microseconds (one doubleword in each logical memory).

Comprehensive error checking has been incorporated in the basic design of the channel hardware. Checking is performed on the control logic in most areas and standard parity checking is done on the data flow. Improved error recovery hardware has also been included (discussed fully in the RAS section).

The Channel-to-Channel Adapter feature available for the Model 65, which permits two System/360 channels to be interconnected, is also an optional feature for the Model 165. The Channel-to-Channel Adapter itself cannot be installed on a 2880 or a 2870 channel; however, a 2880 or a 2870 can be connected to an adapter installed on another channel. Thus, a 2880 or a 2870 can be interconnected to a 2860 channel, a Model 155 channel, a Model 50 channel, etc., that has the Channel-to-Channel Adapter attached.

THE 2880 BLOCK MULTIPLEXER CHANNEL

The 2880 can operate either as a selector or a block multiplexer channel. The setting of a channel mode bit in a control register determines the mode in which 2880 channels operate. The mode bit is set to selector mode at IPL and on system reset and can be altered by programming at any time. When a START I/O instruction is issued to a 2880, the setting of the channel mode bit determines the mode in which the addressed subchannel will operate.

The new START I/O FAST RELEASE instruction can be used with 2880 Block Multiplexer channels. This instruction differs from a START I/O in that START I/O FAST RELEASE permits the CPU to execute the next sequential instruction sooner. That is, if the addressed 2880 is not busy when START I/O FAST RELEASE is issued, the assumption is made that the I/O operation can be started and the CPU is not held up awaiting a response from the device's control unit. If it is determined

later that the I/O operation could not be started, an I/O interrupt occurs to indicate this fact.

Block multiplexer mode is designed to increase system throughput by increasing the amount of data entering and leaving the system in a given period of time (the effective data rate). Better use of 2880 channel time is achieved by operating the channel in block multiplexing mode. A single 2880 block multiplexer channel can support interleaved, concurrent execution of multiple high-speed channel programs. The block multiplexer channel can be shared by multiple high-speed I/O devices operating concurrently, just as the byte multiplexer can be shared by multiple low-speed devices.

Like the 2870 Multiplexer, the 2880 Block Multiplexer Channel has multiple subchannels, each of which has an associated UCW and can support one I/O operation. A 2880 channel has either 64 nonshared subchannels or one shared and 56 nonshared subchannels.

A block multiplexer channel functions differently from a selector channel in the way in which it handles command-chained channel programs. A selector channel executing a command-chained channel program is busy during the entire time the channel program is in operation, whether data transfer is occurring or not. A block multiplexer channel executing a command-chained channel program has the ability to disconnect from the operational channel program during certain non-data transfer operations. That is, a block multiplexer channel can be freed during a nonproductive activity, for example, during disk seeking and most record positioning, thereby allowing more data to be transferred per unit of channel busy time.

Block multiplexing operates as follows. Assume a block multiplexer channel is executing a channel program consisting of multiple command-chained CCW's. When channel end is presented without concurrent device end, the channel disconnects from the channel program and becomes available for an I/O operation on another device, even though the disconnected channel program is not complete. At channel disconnect time the subchannel and the device's control unit retain the information necessary to restart the disconnected channel program.

When the device signals that it is again ready for the channel (by presenting device end), its control unit attempts to regain use of the channel. If the channel is free at this time, the active channel registers are reloaded with the information previously saved (in the device's UCW) and the disconnected channel program is resumed at the appropriate CCW. If the channel is busy when reconnection is requested, the device must wait until it becomes available. Once multiple channel programs have been initiated on one channel, the interleaving of data transfer operations is controlled by block multiplexer channel hardware and the control units of the devices operating in block multiplexing mode.

To facilitate channel scheduling, a new interrupt condition, called channel available, has been defined for block multiplexer channels. At disconnect time for a channel program, the channel is available for the resumption of an uncompleted channel program previously started, or another channel program can be initiated. A channel available interrupt occurs at disconnect time to indicate channel availability if a START I/O, TEST I/O, TEST CHANNEL, or HALT DEVICE instruction was issued previously while the block multiplexer channel was busy.

Two additional facts should be noted about block multiplexer channel operations:

1. When multiple channel programs are operating concurrently in block multiplexing mode, a device can regain control of the

channel only when the channel is not busy. Thus, only cyclic devices (such as direct access devices with rotational position sensing) or buffered devices (such as the 2540 Card Read Punch and the 1403 Printer) can disconnect during the execution of a command-chained channel program on a block multiplexer channel and resume operation later.

2. Data transfer operations for concurrently operating devices on a block multiplexer channel are interleaved on a first-come, first-served basis as the desired records become available. Thus, devices are serviced in the order in which their records become available, not necessarily in the order in which their channel programs are initiated.

10:25 BLOCK MULTIPLEXING OPERATIONS WITH ROTATIONAL POSITION SENSING DEVICES

Rotational position sensing and multiple requesting are standard on 3330 and 2305 facilities. These two functions, together with block multiplexing, are designed to increase system throughput by increasing channel throughput.

The presence of RPS in the control unit of a direct access device enables the device to operate in block multiplexing mode. The use of rotational position sensing reduces the number of channel programs that have to be initiated for direct access devices that require an arm positioning seek (such as the 3330 facility), frees channels more often during direct access device operations - specifically, during most of the time required to position a track to a desired record - and permits disk channel programs to be initiated sooner on block multiplexer channels than is possible with selector channels.

Multiple requesting is implemented in a direct access device control unit to enable it to handle concurrent execution of multiple RPS channel programs. The control unit of the 3330 facility, for example, can simultaneously control eight RPS channel programs, one on each of its drives.

In order to overlap seek operations for current direct access devices without RPS, channel scheduling routines must initiate two channel programs for each record read or write. The first is a stand-alone seek, which frees the channel as soon as the control unit accepts the seek address. (The control unit is also free during arm movement.) At the completion of the seek, a device-end interrupt is presented, and the data transfer channel program is subsequently initiated to search for the desired record and transfer the data. A selector channel is busy during the entire search operation that locates the desired disk record on the track. Search time can be significantly greater than data transfer time for disk records smaller than half a track in size. Search time averages one-half of a rotation for a read or write (8.3 ms for a 3330) and requires a full rotation, less record write time, for a write verification chained from a write.

Use of RPS reduces the time the channel is busy searching for a disk record. It permits the SEARCH command to be initiated just before the desired record is to come under the read/write heads, that is, when the desired rotational position is reached. To accomplish this, a "sector" concept is employed. The tracks in each cylinder of a direct access device are considered to consist of equally spaced sectors (the number of sectors varies by device). Track formatting is unchanged but each record has a sector location as well as a record address. A sector is not physically indicated on disk tracks, but is the length of the track arc that passes under the read/write heads in one sector

time. For the 3330 facility, for example, sector time is defined to be approximately 130 microseconds. Thus, there are 128 sectors per track on the 3330.

A disk control unit with RPS and multiple requesting can determine the sector currently under the heads of each of its drives. A sector counter is contained in each drive. It is incremented once every sector time period and set to zero each time the index marker passes under the heads. The sector in which a record falls is a function of the length of all records that precede it and of its sequential position on the track. Therefore, sector location can be calculated for fixed-length records.

Two new disk commands are provided for use with rotational position sensing:

SET SECTOR
READ SECTOR

If the sector address of a record is known or can be calculated, a SET SECTOR command can be included in the disk channel program to cause the control unit to look for the designated sector. Once the control unit accepts the sector number provided by a SET SECTOR command, both the block multiplexer channel and the disk control unit disconnect and are available for another I/O operation. When SET SECTOR is used for positioning, the time the channel is busy searching for a record is reduced from an average of 8.3 ms to an average of 250 microseconds for the 3330 facility. (Allowing for the worst case of speed variation and for disk pack interchange, the search time for a record, from sector found to beginning of desired record, can vary from 120 to 380 microseconds on a 3330 facility.)

The READ SECTOR command is useful for sequential disk processing and for write verification. When chained from a READ, WRITE, or SEARCH command, READ SECTOR provides the sector number required to access the record processed by the previous CCW. This sector number can be used to reposition the track to the record in order to verify the record just written or in order to read or write the next sequential record. These two new sector commands, used in conjunction with the block multiplexer channel, permit a single command-chained channel program to be initiated for each disk operation to free the channel and disk control unit during seek and rotational positioning operations.

When record ID is known, the two channel programs shown below can be used to retrieve a record directly from an OS BDAM data set on a direct access device without RPS, such as the 2314 (key was not written). The seek operation can be overlapped with other seeks and one data transfer operation on the same selector channel.

Channel program 1. Initiate the stand-alone seek to position the disk arm.

Command Chaining Flag	Command		Selector Channel and Disk Control Unit Status
	SEEK	(Seek address)	Free as soon as the control unit accepts the seek address

Channel program 2. Initiate the data transfer operation after the seek is complete.

Command Chaining Flag	Command		Selector Channel and Disk Control Unit Status
CC	SEARCH ID EQ	(ID - sequential position on the track)	Busy (for 12.5 ms on the average for a 2314)
CC	TIC	(Back to search if ID not equal)	
	READ DATA	(Processor storage address of input area)	Busy

When the sector address is known or can be calculated, the channel program below can be used to retrieve a record from the same BDAM data set on a 3330 facility attached to a block multiplexer channel. The records are fixed-length standard format and sector numbers are calculated from record ID (by data management).

Channel program 1. Initiate the seek and data transfer operation.

Command Chaining Flag	Command		Block Multiplexer Channel and Disk Control Unit Status
CC	SEEK	(Seek address)	Free during arm motion
CC	SET SECTOR	(Sector number of sector preceding desired record)	Free until sector found
CC	SEARCH ID EQ	(ID - sequential position on track)	Busy (250 microseconds average for a 3330)
CC	TIC	(Back to search if ID is not equal. With the logic shown, the first ID inspected normally is that of the desired record and the TIC command is not executed.)	Busy
	READ DATA	(Processor storage address of input area)	Busy

The preceding example indicates the advantages of rotational position sensing used with block multiplexing:

- Only one channel program is required to locate a disk record and transfer the data, thereby eliminating a stand-alone-seek I/O interrupt and the I/O supervisor processing required to schedule

a data transfer channel program. A channel available interrupt may occur, however, during channel program execution.

- The channel and disk control unit are free during arm motion and rotational positioning, allowing seek and set sector operations to be overlapped with other I/O operations on that control unit and channel. Implementation of multiple requesting permits a disk control unit to control concurrent execution of multiple RPS channel programs in order to overlap seek and set sector operations for its drives.

Performance improvement gains achieved on block multiplexer channels are not due entirely to the fact that direct access device rotational delays are overlapped. Also important is the ability to initiate seek commands a number of milliseconds earlier because a block multiplexer channel is free. The initiation of stand-alone seeks on a selector channel is delayed during search and data transfer operations. On a block multiplexer channel, seeks can be initiated during rotational positioning, since the channel and disk control unit are not busy.

The concepts of rotational position sensing as described for the 3330 facility apply also to the 2305 facility. Since an arm positioning seek is not required for the 2305, the channel program issued by an RPS access method or a user can be the same as shown for the 3330 facility except that the first command will not be a seek requiring mechanical motion.

SUMMARY OF BLOCK MULTIPLEXING OPERATIONS WITH I/O DEVICES

The following summarizes how direct access devices without and with RPS and other I/O devices operate on a 2880 Block Multiplexer Channel when a command-chained channel program is executed.

1. Direct access devices without RPS (2311, 2321, 2314, 2303, 2301) assigned to a shared subchannel of a 2880 channel operate in the same way whether the channel is in selector or block multiplexer mode. That is, the 2880 and the disk control unit are busy during the entire time a command-chained disk channel program is in operation. (There is no disconnection after a chained seek.)

2. The 3330 facility executing a command-chained channel program disconnects after the control unit accepts an arm positioning seek that causes arm movement. Reconnection is attempted when the arm reaches its destination and signals device end. The 3330 facility also disconnects when its control unit accepts a SET SECTOR command. When the sector specified arrives under the read/write heads, the control unit attempts to reconnect and resume the CCW chain. If the channel is busy, the control unit repeats the reconnection procedure each time the specified sector position is reached.

3. The 2305 facility executing a command-chained channel program disconnects when the control unit accepts a SET SECTOR command. Reconnection occurs as described for the 3330 facility.

4. All currently announced tape drives operate exactly the same on a 2880 as on a selector channel. That is, the 2880 is busy during the entire time a command-chained channel program is in operation.

5. Buffered card and print devices (or devices operating with buffered control units) disconnect during the mechanical motion of the device. Reconnection occurs later to fill or empty the

associated buffer. For example, a 1403 Printer attached to a 2821 control unit connected to a 2880 channel operating in block multiplexing mode disconnects from the channel during print time and carriage motion. Reconnection occurs when the channel is free, to transfer the data for the next line to the 2821 buffer in burst mode.

6. Any other I/O device that presents channel end without simultaneous device end disconnects from the 2880 when command chaining if it is operating on a nonshared subchannel.

Section 60 discusses planning for the installation of block multiplexer channels and RPS devices.

10:30 STANDARD AND OPTIONAL SYSTEM FEATURES

STANDARD FEATURES

Standard features included on the System/370 Model 165 are:

- Instruction set that includes binary, decimal, and floating-point (including extended precision) arithmetic, the new general purpose instructions, and the instructions required to handle the new architecture. New System/370 instructions are the following:

  COMPARE LOGICAL CHARACTERS UNDER MASK
  COMPARE LOGICAL LONG
  HALT DEVICE
  INSERT CHARACTERS UNDER MASK
  LOAD CONTROL
  MOVE LONG
  SET CLOCK
  SHIFT AND ROUND DECIMAL
  START I/O FAST RELEASE
  STORE CHANNEL ID
  STORE CHARACTERS UNDER MASK
  STORE CLOCK
  STORE CONTROL
  STORE CPU ID
- CPU retry
- Interval timer - 3.33 ms resolution
- Time of day clock
- Expanded machine check interrupt
- Direct Control
- ECC on processor storage
- Byte boundary alignment
- Storage and Fetch Protection
- High-speed buffer storage - 8K
- Attachment feature for up to seven channels (2870, 2860, and 2880)
- Channel retry
- Writable control storage

OPTIONAL FEATURES

The following optional features can be field installed unless otherwise noted:

- High-Speed Multiply (plant installation)
- 709/7090/7094/7094II Compatibility Feature (plant installation)
- 7070/7074 Compatibility Feature (plant installation)
- 7080 Compatibility Feature (plant installation)
- Buffer Expansion feature for inclusion of a 16K buffer
- Channel-to-Channel Adapter - on 2860 only
- Extended Channels (up to 12 addressable channels)
- Remote Operator Console Panel Attachment

Note: Compatibility features are mutually exclusive.

SECTION 20: I/O DEVICES

20:05 I/O DEVICE SUPPORT

Most presently announced I/O devices that can be attached to System/360 Models 65 and above can be attached to the System/370 Model 165. The following I/O devices are not included in a standard Model 165 configuration:

1017	Paper Tape Reader
1018	Paper Tape Punch
1052	Printer-Keyboard, Models 3, 5, 6, and 8
1231	Optical Mark Page Reader
1259	Magnetic Character Reader
1285	Optical Reader
1404	Printer
1412	Magnetic Character Reader
1418	Optical Character Reader
1428	Alphameric Optical Reader
1445	Printer
1827	Data Control Unit (for attachment of 1800 system analog and/or digital control units to the Model 165)
2302	Disk Storage
2415	Magnetic Tape Unit and Control
2671	Paper Tape Reader
7340	Hypertape Drive
7772	Audio Response Unit

In addition, 2361 Core Storage cannot be attached to the Model 165.

New I/O devices for the Model 165 are:

- The 3330 facility - attaches only to a block multiplexer channel
- The 2305 facility Models 1 and 2 - attaches only to a block multiplexer channel
- The 3211 Printer - attaches to any Model 165 channel
- The 3803/3420 Magnetic Tape Subsystem - attaches to any Model 165 channel

The 3330 and the 2305 facilities represent significant advancements in direct access device technology. They provide larger online data capacity and faster data rates and access than heretofore, as well as expanded error correction features. Both have rotational position sensing and multiple requesting as standard features.

The 3330 represents the latest direct access device with removable, interchangeable disk packs. It embodies new data recording and access technology. The 2305 facility is a major extension of the nonremovable, high-speed, fixed-head direct access storage concept.

The 3803/3420 Magnetic Tape Subsystem incorporates all the latest advances in tape speed, design, and technology. It offers new features and enhanced reliability, availability, and serviceability to 2400-series magnetic tape users.

The major new characteristics of the 3330 and 2305 facilities, the 3211 Printer and the 3803/3420 tape subsystem are discussed in the following subsections.

20:10 3330 DISK STORAGE AND 3830 STORAGE CONTROL

The 3330 facility is a modular, large-capacity, high-performance direct access storage subsystem. The 3330 facility consists of 3830 Storage Control and from one to four 3330 Disk Storage modules. A 3330 module contains a pair of independent disk storage drives as shown in Figure 20.10.1. The new, removable 3336 Disk Pack is used for data storage. Usage meters are contained in the 3830 control unit and in each 3330 module.

THIS PAGE LEFT INTENTIONALLY BLANK

SECTION 20: I/O DEVICES

20:05 I/O DEVICE SUPPORT

Most presently announced I/O devices that can be attached to System/360 Models 65 and above can be attached to the System/370 Model 165. The following I/O devices are not included in a standard Model 165 configuration:

1017	Paper Tape Reader
1018	Paper Tape Punch
1052	Printer-Keyboard, Models 3, 5, 6, and 8
1231	Optical Mark Page Reader
1259	Magnetic Character Reader
1285	Optical Reader
1404	Printer
1412	Magnetic Character Reader
1418	Optical Character Reader
1428	Alphameric Optical Reader
1445	Printer
1827	Data Control Unit (for attachment of 1800 system analog and/or digital control units to the Model 165)
2302	Disk Storage
2415	Magnetic Tape Unit and Control
2671	Paper Tape Reader
7340	Hypertape Drive
7772	Audio Response Unit

New I/O devices for the Model 165 are:

- The 3330 facility - attaches only to a block multiplexer channel
- The 2305 facility Models 1 and 2 - attaches only to a block multiplexer channel
- The 3211 Printer - attaches to any Model 165 channel

The 3330 and the 2305 facilities represent significant advancements in direct access device technology. They provide greater online data capacity and faster data rates and access than heretofore, as well as expanded error correction features. Both have rotational position sensing and multiple requesting as standard features.

The 3330 represents the latest generation of direct access device with removable, interchangeable disk packs. The 2305 facility is a major extension of the nonremovable, high-speed, fixed-head direct access storage concept.

The major new characteristics of the 3330 and 2305 facilities and the 3211 Printer are discussed in the following subsections.

20:10 3330 DISK STORAGE AND 3830 STORAGE CONTROL

The 3330 facility is a modular, large-capacity, high-performance direct access storage subsystem. The 3330 facility consists of 3830 Storage Control and from one to four 3330 Disk Storage modules. A 3330 module contains a pair of independent disk storage drives as shown in Figure 20.10.1. The new, removable 3336 Disk Pack is used for data storage.

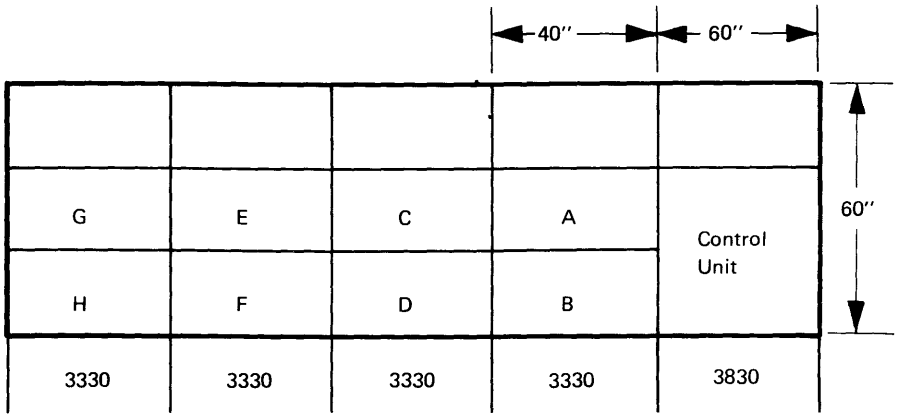

Figure 20.10.1. The 3330 facility

Facility configurations and maximum capacities, using full track records, are shown below:

3830 Storage Control + one 3330 module	200 megabytes
3830 Storage Control + two 3330 modules	400 megabytes
3830 Storage Control + three 3330 modules	600 megabytes
3830 Storage Control + four 3330 modules	800 megabytes

Drives are mounted in powered drawers that are opened and closed by a switch on the operator control panel on the 3330 module. Logical address plugs are supplied, as for the 2314, in addition to a CE service plug. The latter is used when inline diagnostics are to be executed.

Functionally, the 3330 facility provides more capabilities than the 2314, particularly in the areas of performance and availability. The 3330 facility supports all the standard 2314 commands (except the file scan commands) in addition to several new operations, including RPS and error recovery commands. The 3330 facility also is an attractive growth device for the 2321 Data Cell Drive.

The new, removable 3336 Disk Packs are interchangeable across 3330 disk drives but are not interchangeable with the 2316 Disk Packs used on 2314 disk drives (Table 20.10.2 compares disk pack characteristics). Like 2316 packs, 3336 Disk Packs will be initialized in the factory with home addresses and capacity records (R0). Up to 20 defective tracks per pack will be flagged and have alternates assigned. The quick DASDI routine (part of the IEHDASDR utility) currently available for processing 1316 and 2316 packs, will support 3336 packs. Quick DASDI writes the volume label, the VTOC, and, if requested, IPL records but bypasses track analysis. It also determines the number of flagged tracks and places this data in the VTOC.

Table 20.10.1 compares the capacity and timing characteristics of the 3330 facility with those of the 2314 facility and the 2321 Data Cell Drive. The increase in capacity achieved by replacing a 2314 or a 2321 with a 3330 depends upon the block size chosen for the data on the 3330. For example, if the 2314 full track block size of 7294 bytes is maintained for a given data set on the 3330 to avoid programming changes, the 3330 yields a 91% increase in full pack capacity (almost twice the capacity). However, reblocking to a full track on the 3330, 13,030 bytes, yields a 242% full pack capacity increase. If there is not enough processor storage available to allocate I/O areas of 13,030 bytes, lowering the 3330 block size used to half a 3330 track yields a 239% increase in full pack capacity.

If a 2321 is replaced by a 3330, six full track blocks of data from the 2321 (2000 bytes/2321 track) can be placed on each 3330 track if full track blocking is used, for a total of 92,112,000 bytes per 3336 pack (12,000 bytes x 7676 tracks per 3336). Thus, slightly over four 3336 packs provide the capacity equivalent of ten data cells, or a full 2321 drive, if full tracks are used. A full ten data cells, blocked full track, also can be contained in slightly more than four 3336 packs if half-track blocking is used on the 3336.

Table 20.10.1. Capacity and timing characteristics of the 3330 and 2314 facilities and the 2321 Data Cell Drive

Characteristics	3330	2314	2321
Capacity in bytes truncated to the nearest thousand (full track records)			
Pack or cell	100,018,000	29,176,000	39,200,000
Facility or Data Cell Drive			
2 drives/cells	200,036,000	58,352,000	78,400,000
4 drives/cells	400,073,000	116,704,000	156,800,000
6 drives/cells	600,109,000	175,056,000	235,200,000
8 drives/cells	800,146,000	233,408,000	313,600,000
10 cells	-	-	392,000,000
Access time in ms			
Maximum	55	130	600 (for strip select and load)
Average	30	60	175 (minimum for strip select and load)
Average cylinder to cylinder	10	25	95 (on a strip)
Rotation time in ms	16.7	25	50 (strip on drum)
Rotation speed (rpm)	3600	2400	1200
Data transfer rate (KB)	806	312	55

Table 20.10.2. 3336 and 2316 Disk Pack characteristics

Characteristics	3336	2316
Number of disks per pack	12	13
Number of recording disks	10	11
Number of recording surfaces (recorded tracks per pack)	19	20
Disk thickness in inches	.075	.050
Disk diameter in inches	14	14
Disk pack weight in pounds	20	15
Disk pack maximum capacity in millions of bytes	100	29.1
Full track capacity in bytes	13,030	7294
Cylinders per pack	404 plus 7 alternates	200 plus 3 alternates
Tracks per cylinder	19	20
Tracks per pack	7676	4000

Self-formatting records are written on 3336 packs in the same manner as they are on 2316 packs. However, each physical area written (count, key, and data) has a field of error correction code appended to it for the purpose of data validity checking by the control unit instead of the cyclic check area used on the 2314.

The 3830 control unit is microprogram controlled. Read/write monolithic storage contained in the control unit is used for microprogram residence. The control unit also contains a device that reads interchangeable magnetic cartridges (like the device in the system console unit). This device is used for microprogram backup storage and for storage of nonresident diagnostics for the 3330 facility. During a 3330 facility power-on sequence, the functional microprogram is loaded from the device into control storage within the control unit. Therefore, many engineering changes can be installed merely by replacing the magnetic disk cartridge in use with another cartridge that contains the new microprogram.

The 3330 facility also incorporates new error detection, correction, and logging features, designed to improve its availability and serviceability. The following features are implemented in the 3330 that are not provided by previously announced direct access devices:

- I/O error routine correction of recoverable data errors on read operations with data supplied by the control unit in sense bytes
- Command retry initiated by the control unit to attempt hardware correction of certain errors without programming assistance
- Error logging by the control unit in its control storage of successful command retry operations
- Inline diagnostic tests, contained on magnetic disk cartridges, which can be run on a single drive to diagnose hardware malfunctions while other drives in the facility continue normal operations. (Inline diagnostics are provided currently only for 2314 facilities.)

Recovery of correctable data area errors. When the control unit detects a correctable data error during the reading of the data portion of a record, it generates the information necessary to correct the erroneous bytes. The sense bytes presented by the control unit contain a pattern of corrective bits and a displacement value to indicate which of the bytes transferred to processor storage contain the errors. The disk error recovery program need only EXCLUSIVE OR (logical operation) the corrective bit pattern with the error bytes in the input area in processor storage to correct the errors.

Command retry. Error correction (without programming assistance) is performed by a channel/control unit command retry procedure without an intervening I/O interrupt in the following five situations:

1. When a correctable data error occurs during a search or read operation on home address, record count, or record key.

   During a search or read operation, the home address, count, or key read from the disk track is placed in a buffer in control storage within the control unit. When a correctable data error occurs, the control unit corrects the data in the buffer and reissues the command that caused the error. During reorientation to the record, the control unit disconnects and frees the block multiplexer channel. When the failing search or read command is reexecuted, the corrected data in the buffer is used instead of the data actually on the track.

2. When an uncorrectable data error is detected on any portion of the record during a read or a search operation.

   The failing CCW is reissued by the control unit twice. If one of the two retries is successful, the channel program continues normally.

3. When a seek malfunction is detected.

   The control unit retries the command ten times in an attempt to position the arm correctly.

4. When an alternate or defective track condition is recognized before data transfer begins.

   The control unit determines the location of the alternate or defective track (from R0 on the track), initiates a seek to this track, and orients to the index point. When this sequence completes, the original command is reissued by the control unit. This is a programmed procedure for previously announced System/360 direct access devices.

5. When a command overrun (or late command chaining) condition occurs during execution of a channel program because of interference from another channel or the CPU.

   The control unit initiates a retry of the command that was late.

Error logging. Usage and error counters for each drive in the facility are maintained continuously in the control unit. The usage counters are used to accumulate the number of bytes read and seeks issued. The error counters are used to accumulate the number of seek, correctable data, and uncorrectable data errors that were retried successfully by a command retry procedure, as already described. When a counter reaches its threshold or when a pack is removed from a drive, the control unit indicates the condition via a unit check when the next I/O operation is initiated to the drive. Counter data can be obtained and counters can be reset by issuing a SENSE or READ LOG command. These statistics can then be logged in the system error data set for later diagnosis.

Inline diagnostic tests. The 3830 control unit can execute diagnostic tests on a malfunctioning drive while normal operations take place on the remaining drives in the facility. When the CE inserts the service address plug in the malfunctioning drive, diagnostic programs contained on a magnetic disk cartridge are read by the device in the control unit. Diagnostics can be executed on that drive by the customer engineer using the CE panel on the 3830 control unit. Operationally, the drive is offline to the control unit, and physically, the drive is offline to the operating system.

Online testing of the 3330 facility can be performed under OLTEP control as usual. Both OLT's and diagnostic programs contained in the OLT library can be executed on a malfunctioning drive via OLTEP. The diagnostic tests are loaded into control storage in the control unit from the OLT library. Operationally, the 3330 drive is online to the control unit but is logically offline to the operating system.

Inline and online testing allows CE diagnosis and repair of most 3330 failures without the necessity of taking the entire 3330 facility out of the system configuration.

The 3330 facility offers more than additional capacity, fast access, and attractive price performance. The 3330 facility is actually a subsystem in itself. The control unit can control the concurrent execution of one RPS channel program on each of its drives and can handle functions, such as error correction and logging, that normally must be programmed, thereby relieving the control program of these activities. In addition, the availability and serviceability of the 3330 are improved by the implementation of new automatic error correction features, by use of inline diagnostics, and by the speed

and ease of engineering change installation. These factors add to the improvement of total system availability.

Table 20.10.3. Hardware features of 3330 and 2314 facilities

Feature	3330	2314
Number of drives per facility	2, 4, 6, or 8	1, 2, 3, 4, 5, 6, 7, or 8 (A ninth can be included as a spare only.)
Removable interchangeable disk packs	Yes	Yes
Removable address plugs	Yes	Yes
Record overflow feature	Standard	Standard
File scan feature	Not available	Standard
Multiple track operations	Standard	Standard
Two-channel switch	Optional	Optional
Second control unit (to permit two concurrent data transfer operations on a facility)	Not available	Optional (2844 Auxiliary Storage Control)
Rotational position sensing	Standard (128 sectors/track)	Not available
Multiple requesting	The control unit can concurrently handle one channel program on each of its drives.	Not available
Command retry by control unit and channel	Standard	Not implemented
Error correction data presented by control unit	Yes	No
Writable storage in control unit loaded from a magnetic disk cartridge	Yes	No
Inline diagnostic tests initiated via the CE panel in the facility	Standard	Standard
Inline diagnostic tests initiated via the system console	Standard	Not implemented

20:15 THE 2305 FIXED HEAD STORAGE MODULE AND 2835 STORAGE CONTROL MODELS 1 AND 2

One or two 2305 Fixed Head Storage Modules can be attached to a 2835 Storage Control unit. Each module contains six nonremovable rotating disks on which data is recorded. Read/write heads, called recording elements, are fixed in position to access each track on the twelve recording surfaces so that no arm motion is required. (See Table 20.15.3 at the end of this subsection for a comparison of the characteristics of the 2305, Models 1 and 2, with those of 2301 Drum Storage.)

Spare, or alternate, tracks are provided in 2305 modules and must be wired in by a customer engineer to replace defective recording tracks. However, one spare track is available for assignment by the alternate track assignment utility program when a permanent track error occurs on a recording track during processing. Once a spare has been assigned as an alternate track, the hardware automatically accesses

the alternate track when the defective recording track is addressed. This is called alternate track sparing. Switching to an assigned alternate track must be a programmed action for currently announced System/360 direct access devices.

The 2835 control unit contains new error correction facilities similar to those of the 3830 control unit. Recorded data areas within self-formatting records have correction code bytes (16 for the Model 1 and 10 for the Model 2) appended to them instead of a two-byte cyclic check code. When certain types of data errors occur during the reading of the data portion of a record, the control unit can determine the bits in error and generate correction data. This recovery information is presented to the error routine via the sense bytes and can be used to correct the invalid record in processor storage (as described for the 3330 facility).

A command retry feature is implemented in the 2835 that permits certain types of failing commands to be reissued automatically by the channel, when requested by the control unit, without an intervening I/O interrupt. For example, when a count or key area is read erroneously, the control unit retries the command once. If the error is not corrected by the retry, the control unit corrects the data in its own buffer, reexecutes the failing read, and presents the corrected data from the buffer instead of reading it from the track.

Like the 3830 control unit, the 2835 contains a device that reads magnetic disk cartridges containing the control unit microprogram and diagnostic routines.

DATA RECORDING ON THE MODEL 2

Data tracks on the Model 2 are formatted in the same manner as on System/360 direct access devices except for the absence of a home address on each track. There are 768 recording tracks and 96 spare tracks in one module. One recording element is positioned over each track. Each of the twelve surfaces contains 72 tracks, 64 recording and 8 spare. The spare tracks are interspersed among the 72 tracks so that every ninth track is a spare. Data is recorded serially by bit on each track.

Four nonmovable access mechanisms are positioned around the rotating disks, as shown in Figure 20.15.1. Each access mechanism contains two groups of 9 recording elements per surface (for a total of 16 recording and 2 spare elements) and accesses one-quarter of the tracks on each surface. A group of 8 recording elements accesses every other track. The outermost element group of the access mechanism at the top of Figure 20.15.1 accesses data tracks 1, 3, 5, ..., 15, while data tracks 2, 4, 6, ..., 16 are accessed by the outermost element group of the access mechanism at the bottom of the diagram.

There are 90 sectors per track on the 2305 Model 1 and 180 sectors per track on the Model 2. When RPS is used, search time, from sector found to beginning of desired record, ranges from a minimum of 114 microseconds to a maximum of 198 microseconds.

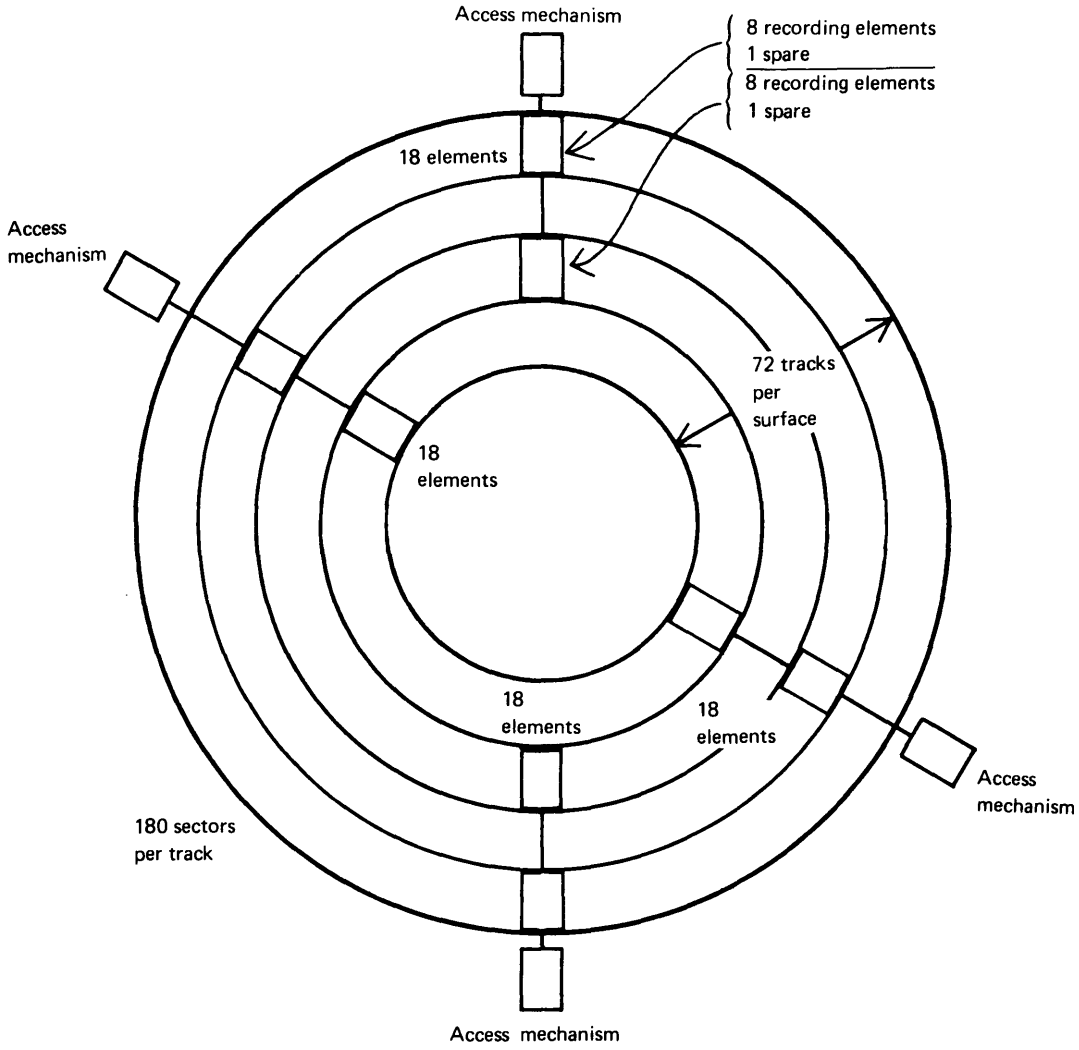

Figure 20.15.1. Top view of a 2305 Model 2 disk surface

DATA RECORDING ON THE MODEL 1

Self-formatting records are written on a Model 1 as on a Model 2. However, data is recorded on the disks of a Model 1 module in a unique way in order to achieve faster data access. The Model 1 module contains the same number of recording elements as the Model 2, but they are positioned in such a way as to double the data transfer rate and halve the average access time. The capacity of the Model 1 is less than half that of the Model 2 because (1) only 384 recording tracks and 48 spare tracks are used, (2) there are two R0 records present on each physical track (360-degree arc) of the Model 1, and (3) the gaps between data areas on a Model 1 track are larger than those of the Model 2 because of the higher data rate of the Model 1.

An addressable recording track on a Model 1 occupies a 180-degree arc on a disk surface rather than a 360-degree arc as on a Model 2.

A recording track in a Model 1 consists of two logical track segments. One track segment occupies a 180-degree arc on the top surface of a disk, while the other track segment occupies a 180-degree arc on the bottom surface of the same disk, directly underneath the top segment. Two recording elements are paired to access each addressable track in parallel. One element records data on the track segment on the top of the disk surface while the other element records data on the track segment on the bottom of the disk surface. Data is recorded in parallel, two bits at a time, one bit on each segment. In addition, all odd bytes are recorded on the top track segment, while all even bytes are recorded on the bottom track segment. Half a rotation is required to record a full track of data. All odd track addresses occupy the same half of a disk, while all even track addresses occupy the remaining half. There are 64 recording and 8 spare tracks on each of the six disks in a module.

Conceptually, the 2305 Model 1 can be thought to contain two access mechanisms 180 degrees apart, each of which has 36 pairs of recording elements per disk, as shown in Figure 20.15.2. (In reality there are two pairs of access mechanisms.) Each of the addressable tracks in one module (384 recording and 48 spare) can be accessed by both mechanisms. The 2835 control unit dynamically selects the mechanism closest to the desired record when a command is issued. The maximum rotational delay is, therefore, half a track rotation (5 ms) and the average delay is 2.5 ms. (See Figure 20.15.2.)

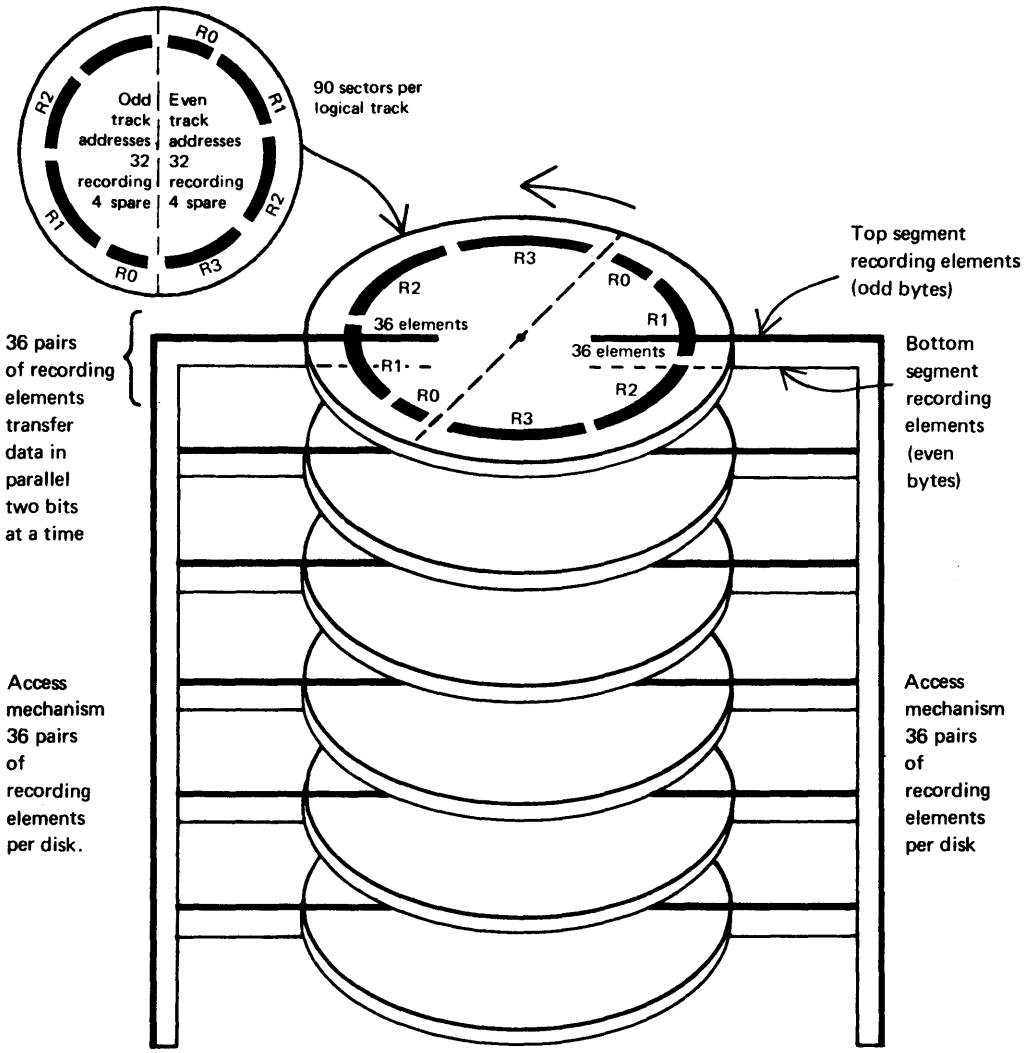

64 recording tracks/disk x 6 = 384 recording tracks/module
8 spare tracks/disk x 6 = 48 spare tracks/module
216 pairs of recording elements (read/write heads)/access mechanism

Figure 20.15.2. 2305 Model 1 module

EFFECTIVE CAPACITY OF 2305 FACILITIES

For certain block sizes the effective capacity of a 2305 Model 1 module will be less than that of a 2301 drum. Tables 20.15.1 and 20.15.2 include some of the block sizes for which this is true. Therefore, block size should be chosen carefully.

Table 20.15.1. Effective capacity of the 2305 Model 2, the 2305 Model 1, and the 2301 for various block sizes (DL column) with a 25-byte key

KL = 25 DL in Bytes	Effective Capacity in Bytes		
	2305 Model 2	2305 Model 1	2301
100	2,688,000	729,600	1,320,000
200	4,300,800	1,228,800	2,000,000
300	5,529,600	1,728,000	2,400,000
400	6,144,000	1,996,800	2,640,000
500	6,912,000	2,304,000	2,800,000
600	7,372,800	2,534,400	3,000,000
700	7,526,400	2,688,000	3,080,000
800	7,987,200	2,764,800	3,200,000
900	8,294,400	3,110,400	3,240,000
1000	8,448,000	3,072,000	3,400,000
2000	9,216,000	3,840,000	3,600,000
3000	9,216,000	3,456,000	3,600,000
4000	9,216,000	4,608,000	3,200,000
5000	7,680,000	3,840,000	3,000,000
6000	9,216,000	4,608,000	3,600,000
7000	10,752,000	2,688,000	2,800,000
8000	6,144,000	3,072,000	3,200,000
9000	6,912,000	3,456,000	3,600,000
10,000	7,680,000	3,840,000	4,000,000

Table 20.15.2. Effective capacity of the 2305 Model 2, the 2305 Model 1, and the 2301 for various block sizes (DL column) when records are written without a key

KL = 0 DL in Bytes	Effective Capacity in Bytes		
	2305 Model 2	2305 Model 1	2301
100	3,763,200	1,036,800	1,760,000
200	5,683,200	1,766,400	2,440,000
300	6,681,600	2,188,800	2,820,000
400	7,372,800	2,611,200	3,040,000
500	8,064,000	2,880,000	3,200,000
600	8,294,400	3,225,600	3,360,000
700	8,601,600	3,225,600	3,360,000
800	8,601,600	3,379,200	3,520,000
900	8,985,600	3,456,000	3,420,000
1000	9,216,000	3,840,000	3,600,000
2000	9,216,000	3,840,000	3,600,000
3000	9,216,000	4,608,000	3,600,000
4000	9,216,000	4,608,000	3,200,000
5000	7,680,000	3,840,000	4,000,000
6000	9,216,000	4,608,000	3,600,000
7000	10,752,000	2,688,000	2,800,000
8000	6,144,000	3,072,000	3,200,000
9000	6,912,000	3,456,000	3,600,000
10,000	7,680,000	3,840,000	4,000,000

ROTATIONAL POSITION SENSING AND MULTIPLE REQUESTING

RPS is a standard feature of the 2835 control unit, as is the other new capability called multiple requesting, which allows up to eight channel programs to be active concurrently on each of the two 2305 modules that can be attached to the control unit. In other words, a 2305 module can be viewed as eight logical devices, although physically it is only one device.

As described previously, rotational position sensing and block multiplexing permit a direct access device to disconnect during SET SECTOR operations. These facilities, used in conjunction with the multiple requesting feature, permit concurrent operations to take place on each 2305 module in a facility. Thus, the effective data rate of a module can be increased substantially.

The multiple requesting capability is implemented by associating eight logical device addresses (0-7) with a 2305 module. Each logical device address also is assigned to a specific subchannel of the block multiplexer channel and a specific register (0-7) in the 2835 control unit. When a channel program is initiated, it is associated with an available logical device address (by the I/O supervisor). (Logical device addresses are not permanently assigned to specific tracks or data sets in a module.) When the SET SECTOR command is issued, the specified sector number is stored in the register in the control unit that is assigned to the logical device address being used for its channel program. Then the control unit disconnects from the channel.

At this point, another channel program with a SET SECTOR command can be accepted by the channel and control unit (assuming neither is busy). This channel program will be initiated using another available logical device address and its assigned control unit register. This process can be repeated for up to eight SET SECTOR commands, so that eight channel programs can be executing concurrently per 2305 module.

Whenever the control unit is not executing a command or is not otherwise busy, it monitors the rotational position counter in the module that is being incremented each sector time period. When the sector number in the counter of a module compares equal with the sector number stored in one of the registers in the control unit and the channel is free, the control unit reconnects and resumes execution of the suspended channel program associated with the logical device address assigned to the control unit register that compared equal (see Figure 20.15.3).

It should be noted that one 2305 module requires eight logical device addresses, each of which requires one subchannel of the 64 available with a 2880 Block Multiplexer Channel. Since each 2835 control unit can have two modules, one 2305 facility uses 16 device addresses and 16 nonshared subchannels of a 2880 Block Multiplexer Channel. In addition, the 2305 Model 1 requires the presence of the Two-Byte Interface feature on the 2880 channel to which it is attached.

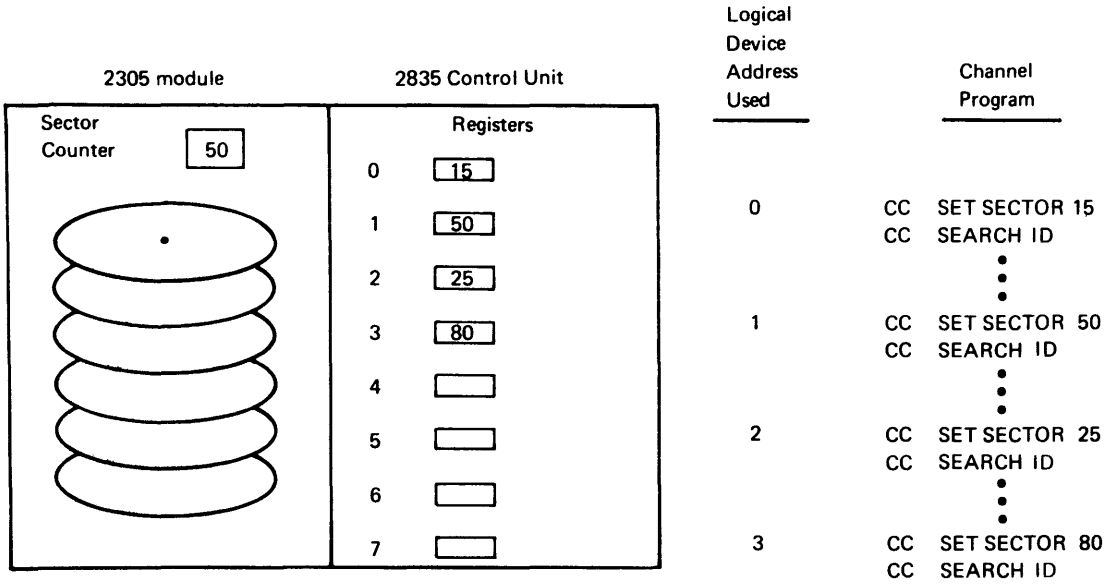

Figure 20.15.3. Multiple requesting on the 2305 facility

Table 20.15.3. 2305 facilities and 2301 Drum Storage characteristics

Characteristic	2305 Model 1	2305 Model 2	2301 Drum
Device type	Six rotating disks with twelve recording surfaces	Same as Model 1	Rotating drum
Module capacity in bytes (full track records, no key)	5,428,224	11,258,880	4,096,600
Number of recording tracks	384	768	200
Number of bytes per track (R0, R1 written without key)	14,136 (Home address is never written on a track)	14,660 (Home address is never written on a track)	20,483 (Home address is always written on each track)
Number of read/write heads (recording elements) per module	864 (768 positioned to access 384 recording tracks. Two elements per data track handle data in parallel. 96 positioned to access 48 spare tracks.)	864 (One positioned to access each of 768 recording and 96 spare tracks.)	880 (One set of four positioned to handle each of 200 recording and 20 alternate tracks. Data is transferred four bits at a time.)
Rotation time (ms)	10	10	17.5
Access time (ms)			
Maximum	5.1	10.25	17.5
Average	2.5	5.0	8.6
Time channel busy searching when SET SECTOR is used (ms)	.114 minimum .198 maximum	.114 minimum .198 maximum	
Rotation speed (rpm)	6000	6000	3400
Data transfer rate in MB	3.0	1.5	1.2
Data validity checking	16 correction code bytes (CCB) appended to each area written	10 correction code bytes appended to each area written	2-byte cyclic check (CC) code appended to each area written
Error recovery performed by the control unit	1. Command retry to retry a failing command without an I/O interrupt	Same as Model 1	Not provided

Table 20.15.3 2305 facilities and 2301 Drum Storage characteristics (continued)

Characteristic	2305 Model 1	2305 Model 2	2301 Drum
	2. Correction of data errors that occur in the data area of a record possible by programming using information in sense bytes		
Rotational position sensing	Standard feature (90 sectors per logical track)	Standard feature (180 sectors per track)	Not available
Multiple requesting (allows concurrent I/O operations on one module)	Standard feature (up to 8 concurrent operations per 2305 module; up to 7 supported by data management)	Standard feature, same as Model 1	Not available, one I/O operation at a time supported
Record Overflow	Standard feature	Standard feature	Standard feature
Two-Channel Switch	Special feature	Special feature	Special feature
A device contained in the control unit	Yes to read the control unit microprogram and diagnostic programs from magnetic disk cartridges	Same as Model 1	No

20:20 THE 3211 PRINTER

The 3211 is a high-speed line printer with front printing and new features designed to reduce operator intervention. The 3211 can print 2000 alphameric lines per minute (with a 48-character set) and is designed for industries and applications, such as utilities, aerospace, finance, communications, and fabrication and assembly, that require high print speeds.

The 3211 has a standard 132-print-position line, which can be expanded to 150 positions as an option. The number of print positions does not affect printing speed. The Universal Character Set (UCS) feature is standard and the interchangeable train cartridge contains 432 graphics. The cartridge character arrangement is unrestricted and can be alphabetic, numeric, or special characters in any combination. When the character arrangement is optimized for specific printing loads, speeds of up to 2500 lines per minute can be attained.

The 3211 attaches to a 3811 Control unit. Unlike some models of the 2821 control unit which can handle multiple devices, a 3811 controls only one 3211 Printer.

New features of the 3211 include a powered forms stacker, an automatic platen, and a tapeless carriage. The powered stacker

mechanism is self-adjusting and automatically rises in increments as the stack of paper mounts. This insures that the stacker mechanism is always the same distance above the top of the stack of forms. The rate of rise during each increment is determined by the setting of the stacker rate knob, which can be adjusted by the operator to produce the best condition for the thickness of the forms being stacked. The stacker also can be raised or lowered manually.

When forms are inserted, the printer platen automatically positions itself close to the train cartridge in accordance with the thickness of the forms. Thus, correct clearance between the platen and the cartridge is achieved without operator intervention. Because of its automatic forms thickness sensing, the 3211 is sensitive to forms with a different degree of thickness at each edge. (For forms limitations, see Form-Design Considerations-System Printers, GA24-3488.)

Forms control paper carriage tape loading and unloading by the operator is eliminated by implementation of a tapeless carriage feature for the 3211. Forms spacing and skipping are controlled by a program-loaded forms control buffer (FCB) contained in the 3811 control unit.

The FCB contains 180 storage positions, each of which corresponds to a print line, that is, a single space of the carriage. Thus, forms up to 22.5 inches in length can be accommodated at 8 lines per inch spacing (or 24 inches at 6 lines per inch). Up to twelve channel codes (1-12), corresponding to the twelve channel positions of the paper carriage tape used on a 1403 Printer, can be stored in the appropriate buffer line positions to control carriage skipping. The FCB can be considered to contain a storage image of a carriage control tape.

A carriage control address register is used to address the FCB and maintain correct line position with respect to the form. This register is incremented as space and skip commands are issued that cause the form to advance. When a SKIP TO CHANNEL command is executed, the carriage control address register is incremented and the form moves until the channel specified is sensed in a line position in buffer storage. If the requested channel number is not found in the buffer, forms movement stops after address position 1 (line 1) has been sensed twice. This prevents runaway forms skipping.

A flag in a buffer storage line position is used to indicate the last line of the form for forms shorter than 180 lines. A flag bit is also used in the first buffer storage position to indicate six or eight lines per inch spacing. The FCB is loaded with the desired forms spacing characters via a LOAD FCB command issued by a program. An error indication is given if an end-of-page flag is not present or if an invalid carriage code is loaded.

Serviceability features, in addition to those provided for the 1403 Printer, are incorporated into the design of the 3211. The fact that a 3811 control unit controls only one 3211 Printer, instead of multiple devices, permits offline repair of the malfunctioning printer or control unit only, without the necessity of removing other operational units from the system.

The 3811 control unit presents six bytes of sense information to identify printer and control unit malfunctions instead of only one, as is provided for the 1403. Certain errors (such as a parity check in the print line buffer) that might be corrected by programmed retry of the print operation are identified in the sense bytes, and carriage motion is suppressed. This permits error recovery without operator intervention if the retry is successful. The additional status data presented can be stored for later analysis and should help speed the diagnosis of hardware malfunctions.

20:25 THE 3803/3420 MAGNETIC TAPE SUBSYSTEM

The new 3803/3420 Magnetic Tape Subsystem consists of 3803 Tape Control and a family of three 3420 Magnetic Tape Units which read and write nine-track, 1600-BPI, phase-encoded, half-inch magnetic tape. The three tape units, Models 3, 5, and 7, have a data rate of 120 KB, 200 KB, and 320 KB, respectively, and up to eight tape units, in any mixture of models, can be attached to a 3803 control unit. This tape subsystem, which embodies a completely new control unit technology, offers price performance improvements, compatibility with existing seven- and nine-track tape volumes and programs, enhanced reliability, availability, and serviceability features, lower cost tape switching capabilities, and standard automated tape-handling features presently available only on 2420 Magnetic Tape Units. (Table 20.25.2 at the end of this subsection compares 3420, 2420, and 2401 tape unit characteristics.)

The 3803/3420 subsystem can be attached to all System/370 models and to System/360 Models 30 to 195 (Model 67 in 65 mode only) except Model 44 for which there is no program support. The tape commands, status responses, and basic sense data of this tape subsystem are compatible with those of 2400-series tape units. Thus, any correctly written, non-time-dependent System/360 program for 2400-series tape units will operate without change on the Model 165 (subject to restrictions stated in Section 10:05) to handle operations on 3803/3420 subsystems with equivalent features installed. That is, existing nine-track 1600-BPI phase-encoded (PE), nine-track 800-BPI non-return-to-zero (NRZI), and seven-track 556/800-BPI NRZI-encoded tapes can be processed on 3420 tape units using existing programs without change to the tape volumes or programs.

The 3803/3420 tape subsystem offers 2400-series tape users the advantages of the latest significant advances in tape speed and design while maintaining media compatibility with existing tape volumes and providing enhanced RAS features. Specifically, the following are provided:

- Data rates of 120 KB, 200 KB, and 320 KB at 1600-BPI density
- Phase-encoded data recording that automatically detects and corrects single-bit read errors in-flight
- A tape transport design that minimizes tape wear and increases reliability, a single-drive capstan to control tape movement that provides faster data access times and rewinds, and more precise control of motor speed to help minimize damage to tape media
- Cartridge loading of tape, automatic tape threading, and a new automatic tape reel hub latch, all to reduce tape setup time
- Dual Density and Seven-Track (mutually exclusive) features to enable a 3420 tape unit to handle either nine-track 800-BPI NRZI and 1600 BPI PE tape or seven-track 556/800-BPI NRZI (BCD or binary) tape
- Flexible, lower cost tape switching implemented in a new compact physical design - a two-channel switch is available also
- Features such as new technology to improve subsystem reliability and new diagnostic facilities to aid serviceability and thereby increase subsystem availability

Phase-encoded recording. The phase-encoded (PE) recording technique offers superior error detection and reliability as compared with the non-return-to-zero (NRZI) technique. In both cases, magnetic recording of one and zero bits is accomplished by means of flux reversals or changes in polarity. In NRZI recording, only one bits are recorded as magnetized spots, and a flux reversal occurs only for one bits. In PE recording both zero and one bits are recorded (the zero bit and one bit being opposite in polarity), and a flux reversal is required in every bit position. Thus, the PE dual flux recording technique differentiates between no recording and the presence of a zero bit, and the absence of any signal is detected as an error.

The positive recording of all zero and one bits in PE eliminates the need for horizontal parity bits (longitudinal redundancy check used in NRZI recording), and vertical parity bits are used to correct single-bit read errors in-flight. During reading, if a single track fails to respond with a suitable pulse in any bit position, reading of the rest of that track is immediately disabled for the remainder of the data block, and the remaining bits for that track are automatically generated by use of the vertical parity bits. In-flight single-track error correction eliminates the time normally lost in backspacing and rereading NRZI tape for correction of single-track dropouts or defects.

Phase encoding offers other advantages. If a string of zeros is recorded on tape, successful reading in NRZI requires close synchronization to "count" the correct number of zeros. With PE, this synchronization is provided by the flux reversal in every bit position; hence, PE recording (and reading) is self-clocking. In addition, each block written on a PE tape is preceded and followed by a coded burst of bits in all tracks to set up the individual track-clocking rates. The read circuitry is designed to recognize these bursts and thereby minimize the effect of noise in the gap.

The critical nature of vertical skew (alignment of bits within a byte) that is imposed by NRZI recording is minimized by this individual track clocking scheme (one clock per track versus one clock for the entire tape subsystem), and by the use of one-byte (nine-bit) capacity skew buffers that can be in the process of collecting up to four data bytes at the same time, as the tape passes the read head. Because of the positive recording of all bits, once a skew buffer contains nine bits, one from each horizontal data track, it is an indication that a byte has been read. Thus, the 3420 can handle the situation in which the tape is not exactly aligned, and bits from up to four adjacent bytes can be read concurrently.

Like 2400-series tape units, the 3420 utilizes a two-gap read/write head that performs readback checking during write operations. The 3420 also has a separate erase head that erases the entire width of the tape during any write operation before waiting occurs. Full-width erasure reduces the likelihood of leaving extraneous bits in interblock gaps or skip areas and minimizes the interchangeability problems that can occur when tape is written on one tape unit and read on another.

Advanced engineering design. The tape path in the 3420 tape unit is designed for "soft handling" of tape volumes to minimize tape wear and thus improve tape reliability. Other features, such as the single-drive capstan and optical tachometers, result in faster data access and rewind times than those of the 2401 and the 2420, for models with comparable data transfer rates.

On a 3420 tape unit, the tape reel is mounted on the right side of the tape transport as on a 2420, instead of on the left as on a 2401 tape unit, so that an inverted tape path exists. As a result, when the tape is loaded in the columns, the recording side touches only the tape cleaner and read/write head. Friction and tape wear are also reduced by the presence of air bearings in the tape transport that provide a thin film of air between the nonrecording surface and each metal bearing.

Use of a single-drive capstan transport for tape movement and optical tachometers for control of motor speed result in several advantages. First, faster access times than those of 2400-series tape unit models are achieved. Access time is defined as the time interval from initiation of a write or forward read command (given when the tape is not at load point) until the first data byte is read or written, assuming the tape is brought up to speed from stopped status. Nominal access times for 3420 Models 3, 5, and 7 are 4.0 ms, 2.9 ms, and 2.0 ms, respectively.

Second, the single-drive capstan can be made to operate faster than normal read/write speed, and in-column rewind is thus implemented. Full reel rewind speeds average 410, 480, and 640 inches per second for Models 3, 5, and 7, respectively. (The 2420 rewinds at about 500 inches per second except for the last 100 feet, which are rewound at the model's rated read speed.) In addition, less time is required to rewind less than a full reel on a 3420 as compared to a 2401 because of faster rewind times achieved by in-column rewinding.

Last, three optical tachometers that monitor motor speed are used to achieve precise control of the speed of both the capstan motor and the tape reel motors. The capstan tachometer measures the size of the interblock gaps (IBG's) created during tape writing. The result is a more consistent IBG size (.6 inches) than is created by 2400-series tape units, which enables more accurate calculation of tape passing time. IBG passing times are 8.0 ms, 4.8 ms, and 3.0 ms for 3420 Models 3, 5, and 7 respectively. These times would be used in calculations for command chained tape operations (reading or writing more than one tape block with a single START I/O instruction). More precise capstan motor speed also results in smoother starts and stops, thereby minimizing tape stretching and breaking.

The two tape reel tachometers measure tape speed as the tape enters and leaves the vacuum columns, and tape speed is adjusted when necessary. The 3420 tape unit is, therefore, less sensitive to voltage changes. More precise control of tape reel motor speed improves rewind speed and minimizes erratic tape stacking during rewinds so that there is less chance of damaging tape edges.

Automatic threading and cartridge loading. These advanced features are standard on all 3420 models and significantly reduce tape mounting and demounting time. Tape threading is automatic for tape reels not enclosed in a wraparound cartridge once the reel (10.5-inch, 8.5-inch, or minireel) is mounted on the tape unit with the tape end placed in the threading chute and the load-rewind button is depressed. The power window is closed, the tape is threaded on the takeup reel, and the tape is loaded in the columns and positioned at load point within ten seconds after the button is depressed for Models 3 and 5. On the Model 7, only seven seconds are required. In addition, unload and rewind/unload operations cause the tape to be completely rewound on the tape reel and the power window to be lowered so that the reel is ready for immediate demounting.

If the tape is enclosed in a wraparound cartridge (10.5-inch reels only), an operator need only mount the cartridge and does not have to place the tape end on the threader chute. Once the load-rewind button is depressed, ten seconds are required to open the cartridge and perform automatic threading. If automatic threading fails on the first try, the 3420 unit automatically rewinds the tape and retries threading. Unload operations rewind and close the cartridge automatically. In addition to fast tape reel mounting, the use of a wraparound cartridge offers other advantages. Handling of the tape reel itself is not required when the tape is used, because the wraparound cartridge is also the shelf storage container. The only time the cartridge need be opened is when it is opened by the 3420 during use. This enhances the reliability of the tape media.

The 3420 tape unit also has a new automatic reel latch instead of the snap type hub latch implemented on newer 2400-series tape units. The operator places the tape reel on the hub and the automatic latch mechanically aligns the reel and then pneumatically locks it in position.

The advantage of these features can be shown by comparing setup times for tape units with and without the autothread feature. A tape study using experienced operators indicated the total time required to remove a tape reel, mount a new reel, thread the tape, and come to ready status was the following:

2401 tape unit - 40 seconds
Autothread tape unit without cartridge - 29 seconds
Autothread tape unit with cartridge - 13 seconds

Single Density, Dual Density, and Seven-Track features. These three features are provided for both 3803 control units and 3420 tape units. Dual Density can be field-installed on a 3803; however, the Seven-Track feature is not recommended for field installation. They are mutally exclusive features. The Dual Density or the Seven-Track NRZI feature can be field installed on a 3420 tape unit only if it is replacing another NRZI feature. (For example, Dual Density can be field installed to replace the Seven-Track but not the Single Density feature.) The Dual Density and Seven-Track features facilitate efficient conversion of existing NRZI-recorded tapes to 1600-BPI phase-encoded format and permit tape volume interchange with other systems that use seven-track 556/800-BPI or nine-track 800-BPI tape. (See Section 60 for a discussion of conversion to 3420 tape units.)

A 3803 control unit with the Single Density feature (without a switching feature) can handle up to eight 3420 tape units (Models 3, 5, and 7) with the companion Single Density feature installed. Only 1600-BPI PE tape can be read and written. When the Dual Density feature is present on the 3803 control unit, both nine-track 1600-BPI PE and nine-track 800-BPI NRZI tape operations can be performed on 3420 units (Models 3, 5, and 7) with the companion Dual Density feature installed. (Tape units with the Single Density feature still handle only nine-track 1600-BPI PE tape.) When the Seven-Track feature is present on the 3803 control unit, seven-track 556/800-BPI NRZI operations (both BCD and binary format) can be performed on 3420 tape units (Models 3, 5, and 7) with the companion Seven-Track feature installed. The data convert and translate facilities are a standard part of the Seven-Track feature. Table 20.25.1 summarizes 3803 control unit capabilities without and with these features.

Tape mode setting is handled as follows. For write operations on nine-track tape units with the Dual Density feature, a mode set command must be issued to establish 1600-BPI PE or 800-BPI NRZI recording mode prior to the first write. Tapes written in PE mode have a format identification burst recorded at load point that differentiates them from NRZI mode tapes. During reading, sensing of this burst automatically puts the tape unit in PE mode. Failure to sense the burst establishes NRZI mode if both the tape unit and control unit have the Dual Density feature. If an attempt is made to read NRZI-mode tape on a unit without the Dual Density feature, an error indication results. Once PE or NRZI mode is established for read operations, it is retained until the tape returns to load point.

For seven-track read and write operations, NRZI mode, density, parity, and use of the data converter or translator are established by issuing a single MODE SET command.

Table 20.25.1. 3803 control unit configurations and capabilities with Single Density, Dual Density, and Seven-Track features

3803 with Single Density Feature	3803 with Dual Density Feature	3803 with Seven-Track Feature (includes data convert and translate)
1. Nine-track, 1600-BPI, PE tape on 3420 Models 3, 5, and 7 with Single Density feature	1. Nine-Track, 1600-BPI, PE tape on 3420 Models 3, 5, and 7 with Single Density feature	1. Nine-Track, 1600-BPI, PE tape on 3420 Models 3, 5, and 7 with Single Density feature
	2. Nine-track, 800-BPI NRZI tapes and nine-track, 1600-BPI PE tapes on 3420 Models 3, 5, 7 with the Dual Density feature.	2. Seven-track, 556/800-BPI, NRZI BCD and binary tapes on 3420 Models 3, 5, and 7 with the Seven-Track feature.
Note: The Single Density, Dual Density, and Seven-Track features are mutually exclusive on the same control unit or the same tape unit.		

Tape-switching features. Tape subsystem configuration flexibility is provided by field installable tape-switching options that permit up to four control units to be switched among up to 16 tape units. While this capability is provided for 2400-series tape units via the 2816 Switching Unit, tape switching for the 3803/3420 subsystem offers the advantages of compact design, reduced cost, and enhanced subsystem availability.

The switching features are built into the 3803 control unit itself so that space for stand-alone switching units is not required. The fact that tape switching features are contained in the 3803 control units being switched (rather than in one unit) also enhances tape subsystem availability. When a switch failure occurs in one control unit, that unit can be switched offline, eliminating the necessity of removing the entire tape switching subsystem from the operative system configuration.

Using combinations of the Communicator and the Two-Control Switch, the Three-Control Switch, or the Four-Control Switch optional features, two, three, or four control units can be configured to be switched among up to 8 or up to 16 tape units. The Communicator must be present in all control units that are to be switched. It allows the control

unit in which it is installed to address tape units that are attached to an interconnected control unit. Figure 20.25.1 shows the switching feature requirements for permissible switching combinations. The switch combinations shown for switching control units among up to 16 tape units are the same that are required for switching control units among up to 8 tape units.

A two control unit switching configuration is required to replace the 2804 and 2404 read-while-write control units. The advantage of the tape switching approach is that for a small price increment better performance is possible. This is true because any two tape operations can be active concurrently in a switched configuration (including two reads or two writes) while the degree of simultaneity achieved using a read-while-write control unit is application dependent. That is, the application must lend itself to reading, then writing (or vice versa).

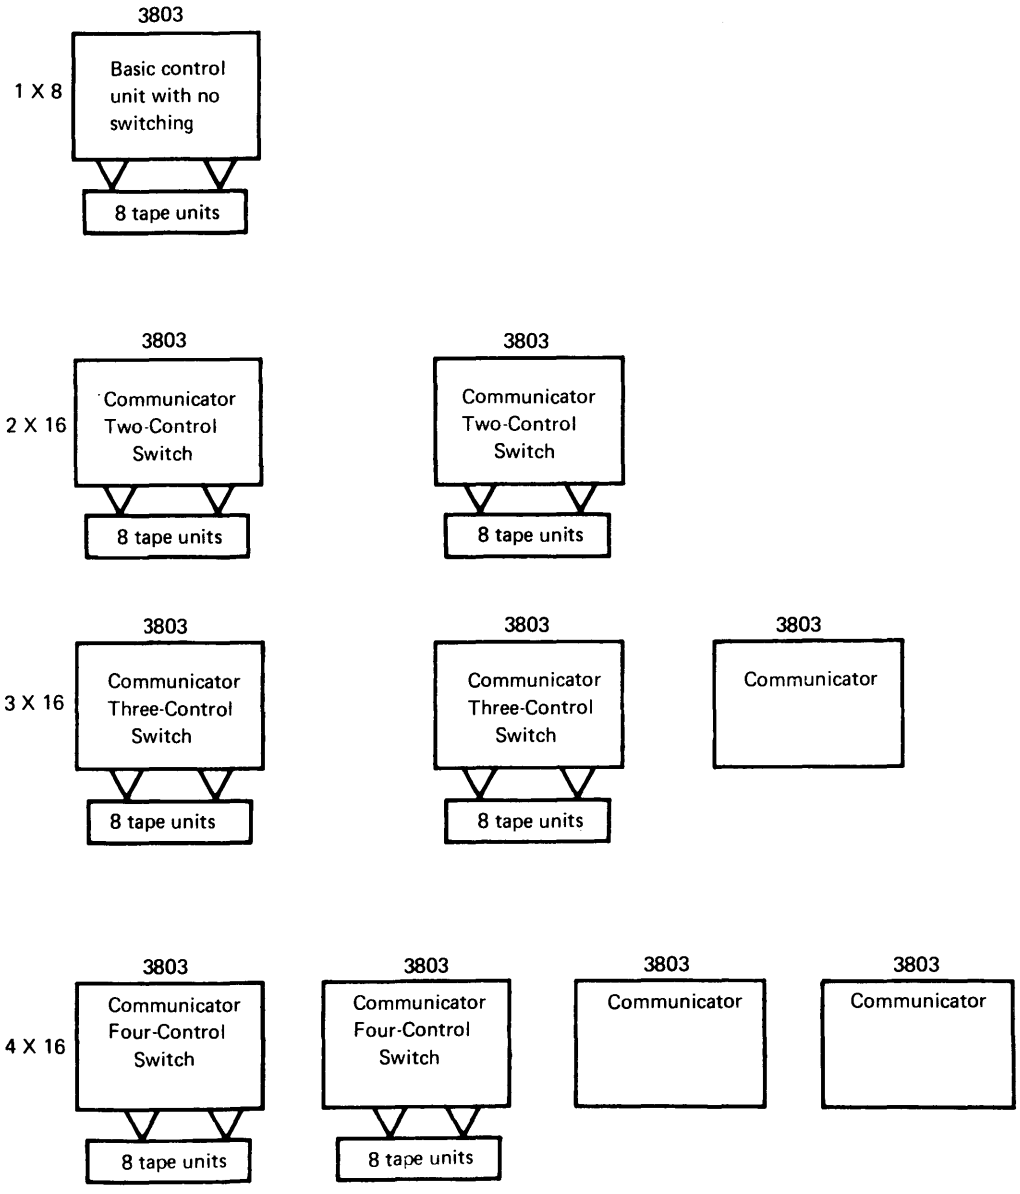

Figure 20.25.1. Tape-switching configurations for the 3803/3420 Magnetic Tape Subsystem

Two-Channel Switch. This optional feature provides switching functions for tape units similar to those provided by the two-channel switch for direct access devices. A 3803 control unit with the two-channel switch installed can be attached to two channels that are in the same system or in two different sytems. Tape units attached to the 3803 can then be accessed via either channel. This feature can be present on a 3803 that also has tape switching features installed.

A 3803 with this switch can be set to allow access to all its tapes by either channel, one channel at a time. If channel A requests an operation when the control unit is busy performing an operation on channel B, then channel A must wait until the control unit becomes available again. If both channels are on the same system, this arrangement essentially provides two channel paths to the tape units on the switched 3803.

A RESERVE CONTROL UNIT command is provided for use with this feature. It permits a channel, via programmming, to maintain exclusive use of the control unit and its tapes until the RELEASE CONTROL UNIT command is issued. These two commands are of benefit when the control unit is shared between two systems.

Tape units on a 3803 with a two-channel switch can also be partitioned. Partitioning is the manual assignment of tape units (via switches) to one channel or the other so that access to each unit is limited to one of the two channels. This facility can be used for backup purposes to switch tapes from one system to another or from one channel to another in the same system.

The two-channel switch for the 3803/3420 subsystem offers configuration flexibility not generally available to 2400-series tape unit users. A two-channel switch currently is provided only for a 2803 Model 1 control unit and can be used only in Model 67 and Model 65 multiprocessing configurations.

Reliability, availability, and serviceability features. The 3803/3420 hardware subsystem has several RAS features, in addition to the reliability and availability features already discussed for the tape media itself.

The 3803 control unit embodies a totally new design. The newest monolithic logic technology is used in the 3803 control unit and, therefore, it offers greater reliability and more compact physical design in comparison to the 2803 control unit. (The 3803 is approximately half the size of a 2803 control unit.) In addition, both logic circuitry and mechanical components in the control and tape units are functionally packaged to enable more rapid fault location and faster replacement.

As a diagnostic aid, additional sense bytes are generated by the microprogram-controlled 3803 control unit. The 3803 uses ROS for microprogram residence. Twenty sense bytes are provided, instead of the six generated by the 2803, certain of which can be used in tracing control unit microprogram malfunctions. Some of the other additional sense bytes identify the control unit and tape unit by serial number, optional features, and engineering change (EC) level.

Two other very significant new serviceability features are microdiagnostics resident in the 3803 control unit and radial attachment of 3420 tape units to the 3803.

Resident microdiagnostics in the 3803 enhance test operations for the 3803/3420 subsystem by relieving the CPU of the execution of most time-dependent tests. Diagnostics in the 3803 are executed via use of a diagnostic command issued by a program.

The 3803 also contains diagnostics that are operative during normal tape processing operations. These diagnostics perform operations such as the monitoring of measurement functions of the tape units. If an irregularity is noted, the control unit generates sense bits to inform the executing program of the malfunction.

Tape subsystem availability is improved by radial attachment of 3420 tape units to the 3803 control unit. That is, each 3420 is cabled directly to the control unit so that any malfunctioning tape unit can be disconnected from the tape subsystem for servicing without disturbing the other tape units. When tape units are attached to the control unit in series (each tape unit cabled to the next tape unit), as are 2400-series units, the entire tape subsystem must be taken offline to uncable a tape unit.

These new features, combined with the use of fewer adjustable parts, are designed to provide optimum tape subsystem availability through better reliability and reduced maintenance time.

In conclusion, the 3803/3420 Magnetic Tape Subsystem offers Model 65 and 75 tape unit users the following:

- Fast data rates, fast access times, and faster rewind time for short files than on 2401 tape units. In-flight correction of single-bit read errors eliminates a backspace and reread procedure and reduces the number of permanent read errors.
- Automatic tape threading and cartridge loading to reduce tape setup time
- 1600-BPI density to reduce the tape volume requirements for multivolume data files recorded at lower densities
- Less tape media wear as a result of the transport design and automatic threading and less tape damage caused by handling if wraparound cartridges are used for tape volume mounting and storage
- Reduced maintenance time because of the transport design (fewer adjustable parts), functional packaging of components, expanded sense bytes, and microdiagnostics resident in the control unit
- Increased tape subsystem availability because of reduced maintenance requirements
- Compatibility with existing 2400-series tape volumes and programs

These advantages, combined with lower subsystem cost and compact, flexible tape-switching capability, make the 3803/3420 Magnetic Tape Subsystem the natural growth path for tape users.

• Table 20.25.2. 3420, 2420, and 2401 Magnetic Tape Unit characteristics

Characteristic	3420 Tape Units			2420 Tape Units		2401 Tape Units			
	Model 3	Model 5	Model 7	Model 5	Model 7	Model 2	Model 3	Model 5	Model 6
Data rate (KB)	120	200	320	160	320	60	90	120	180
Density (bytes/inch)	1600	1600	1600	1600	1600	800	800	1600	1600
Tape speed (inches/sec.)	75	125	200	100	200	75	112.5	75	112.5
Nominal interblock gap size in inches (nine-track)	.6	.6	.6	.6	.6	.6	.6	.6	.6
Nominal read access to data (ms)	4.0	2.9	2.0	3.9	2.5	8	5.3	8	5.3
In-column rewind	Yes	Yes	Yes	Yes	Yes	No	No	No	No
Nominal rewind and unload time (secs.)	76	66	51	78	66	90	66	90	66
Nominal rewind to ready status--full 2400-foot reel (secs.)	70	60	45	72	60	84	60	84	60
Automatic threading	Standard	Standard	Standard	Standard	Standard	Not available	Not available	Not available	Not available
Time to ready status after load button pressed (secs.)	10	10	7	10	7	-	-	-	-
Cartridge Loading (10.5-inch reels only)	Standard	Standard	Standard	Standard	Standard	Not available	Not available	Not available	Not available
Automatic reel latch	Yes	Yes	Yes	No	No	No	No	No	No
Recording technique	PE	PE	PE	PE	PE	NRZI	NRZI	PE	PE

• Table 20.25.2. 3420, 2420, and 2401 Magnetic Tape Unit characteristics (continued)

Characteristic	3420 Tape Units			2420 Tape Units		2401 Tape Units			
	Model 3	Model 5	Model 7	Model 5	Model 7	Model 2	Model 3	Model 5	Model 6
Recording medium (1/2-inch magnetic tape)	IBM Series/ 500 Dynexcel, Heavy Duty, or equivalent 10.5",8.5", 6.5" reels (Use of Mylar* is not recommended.)	Same as Model 3	Same as Model 3	Same as 3420	Same as Model 5	Same as 3420 plus Mylar	Same as Model 2	Same as 3420	Same as 3420
Inverted tape path, single-capstan drive optical tachometers	Yes	Yes	Yes	Yes except for optical tachometers	Same as Model 5	No	No	No	No
Error checking									
Single-track corrections during reading	Automatic	Automatic	Automatic	Automatic	Automatic	Programmed	Programmed	Automatic	Automatic
Vertical redundancy check	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitudinal redundancy check	No	No	No	No	No	Yes	Yes	No	No
Number of sense bytes	20	20	20	6	6	6	6	6	6
Microdiagnostics in control unit	Yes	Yes	Yes	No	No	No	No	No	No
Separate erase head	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seven-Track feature	Optional	Optional	Optional	Not available	Not available	Optional	Optional	Not available	Not available
Densities (BPI)	800	800	800	-	-	800	800	-	-
	556	556	556	-	-	556	556	-	-
	-	-	-	-	-	200	200	-	-
Data rate (KB)									
800 BPI	60	100	160	-	-	60	90	-	-
556 BPI	41.7	69.5	111.2	-	-	41.7	62.5	-	-
200 BPI	-	-	-	-	-	15	22.5	-	-

* Trademark of E.I. Dupont deNemours & Co. (Inc.)

Page of GC20-1730-0
Added 11/20/70
By TNL GN20-2277

• Table 20.25.2. 3420, 2420, and 2401 Magnetic Tape Unit characteristics (continued)

Characteristic	3420 Tape Units			2420 Tape Units		2401 Tape Units			
	Model 3	Model 5	Model 7	Model 5	Model 7	Model 2	Model 3	Model 5	Model 6
Recording technique	NRZI	NRZI	NRZI	-	-	NRZI	NRZI	-	-
IBG size (ins.)	.75	.75	.75	-	-	.75	.75	-	-
Translator	Standard	Standard	Standard	-	-	Standard	Standard	-	-
Data Converter	Standard	Standard	Standard	-	-	Optional	Optional	-	-
Dual Density feature (800/1600 BPI)	Optional	Optional	Optional	Not available	Not available	Not available	Not available	Optional	Optional
Data rate (KB) at 800 BPI	60	100	160	-	-	-	-	60	90
Recording technique at 800 BPI	NRZI	NRZI	NRZI	-	-	-	-	NRZI	NRZI
IBG size at 800 BPI (inches)	.6	.6	.6	-	-	-	-	.6	.6
Control Unit	3803 with optional Seven-Track or Dual Density feature (not both). Read while write (RWW) capability is not provided.	Same as Model 3	Same as Model 3	2803 Model 2. 2420 units can be mixed with 2401 Models 1 to 6 if required attachments are present on 2803	Same as Model 5	2803, 2804 (RWW) Model 1 with optional Seven-Track Compatibility feature 2803, 2804 Model 2 with optional Seven-Track, Nine-Track, or Seven- and Nine-Track Compatibility feature	Same as Model 2	2803, 2804 (RWW) Model 2 with optional Seven-Track, Nine-Track, or Seven- and Nine-Track Compatibility feature	Same as Model 5
Tape switching	2 x 16 3 x 16 4 x 16 (Switching features in 3803 control units)	Same as Model 3	Same as Model 3	2 x 16 3 x 16 4 x 16 (Requires one or two 2816 units)	Same as Model 5	Same as 2420	Same as 2420	Same as 2420	Same as 2420
Two-Channel Switch	Optional	Optional	Optional	Not available	Not available	Optional on 2803 Model 1 for Model 67 and MP65 systems only	Same as Model 2	Not available	Not available

THIS PAGE LEFT INTENTIONALLY BLANK

30:05 TRENDS IN DATA PROCESSING AND PROGRAMMING SYSTEMS

The Model 165 and its programming systems support have been designed to operate in the data processing environment that has been emerging since introduction of System/360.

Significant trends are the following:

- Growth toward more multiprogramming to improve system throughput. Multiprogramming also permits the user to install new applications, such as small teleprocessing inquiry, graphics, or time-sharing applications, that would not otherwise justify a dedicated system. Multiprogramming support also has encouraged the growth of new computer environments, as indicated by the items that follow.

- Growth of integrated emulation, that is, concurrent native and emulation mode processing on one system. The execution of emulators under operating system control improves system throughput because emulators can use control program facilities (stacked job execution, data management functions, etc.) and because native mode and emulator jobs can be scheduled to operate concurrently to utilize available system resources efficiently. The use of integrated emulators eliminates most reprogramming and eases transition from one system to another, permitting the user to expend his efforts extending and adding applications.

- Greater use of high-level languages, such as COBOL, FORTRAN, and PL/I, for applications programming. The cost of programming has been increasing while the cost of computing hardware has been decreasing. More productive use of programmers can be achieved by the use of high-level languages. Improvements to compile times and to the size and execution speed of code produced by high-level language translators have been made and continue to be made. The support of many more functions within high-level languages also has increased their use, and the growth of interactive computing has stimulated the addition of even more facilities.

- Growth of teleprocessing applications such as remote inquiry, message switching, data collection, and management information systems. The ability of System/360 and System/370 to handle teleprocessing and batch processing in one system eliminates the necessity of dedicated special purpose systems.

- Growth of remote computing activities, such as remote job entry and interactive computing (or time-sharing) in both a nondedicated and dedicated environment. Remote computing offers (1) fast turnaround for batch work submitted from remote locations, (2) remote user access to the large computing facilities and data base available at the central installation, and (3) interactive problem solving on a regular or a nonscheduled basis for personnel in locations remote from the central computer. In-house interactive computing is growing also as users attempt to use programmer time more efficiently.

- Growth toward large, online data base systems. The growth in the marketplace of remote computing, time sharing, and real-time applications necessitates the instant availability of more and more data. High-capacity, fast, reliable direct access devices are required for this type of computing environment.

IBM programming systems support of these trends in data processing has been growing and continues to be extended. The System/370 Model 165 offers the hardware, I/O devices, and performance capability required by the expanding computing environment.

30:10 OS SUPPORT

It is expected that the majority of the users who upgrade to the Model 165 will use OS as their operating system. The design of OS and the facilities it currently provides make it a particularly suitable base upon which to build additional support of the data processing trends discussed. Some of the major OS features currently available or announced are the following:

- Priority and job class scheduling for increased resource utilization
- Multiprogramming support of up to 15 user jobs and multitasking facilities
- Automatic data transcription performed by the control program (reader interpreters and output writers) concurrently with user job execution
- Dynamic resource allocation by the control program
- Multiple console support (MCS) and device-independent display operator console support (DIDOCS)
- Extensive teleprocessing and graphics support
- Remote job entry and conversational remote job entry
- Support of a wide range of interactive (time-sharing) environments

OS will be modified and extended in future releases so that it supports Model 165 hardware. Appropriate alteration of the resident portion of a control program (nucleus) generated for a Model 165 will accommodate the fixed storage area of lower storage in the Model 165. OS for the Model 165 includes currently announced OS facilities and contains additional support to handle new Model 165 hardware features and I/O devices. Emphasis has also been placed on extending error recovery, recording, diagnostic, and repair procedures.

OS support of Model 165 features will be provided as follows (Programming systems support of RAS features is discussed in Section 50).

New instructions. The Assembler F (Type I) and Assembler H (program product) language translators will include mnemonics for the general purpose and other new instructions for the Model 165 so that these instructions can be used in user-written Assembler Language programs. The currently available OS high-level language translators will not generate the six new general purpose instructions.

Extended precision floating point. Assemblers F and H will include support of the extended precision floating-point data format and instructions. In addition, extended precision will be supported by the FORTRAN H and the PL/I Optimizing Compiler program products.

The implementation of extended precision support in FORTRAN H and PL/I is such that the language translator and the processing programs generated to include extended precision operations can operate on a System/370 or System/360 with or without extended precision hardware. The language translator contains extended precision instructions and

generates them for processing programs that use extended precision data. A program check interrupt occurs if an extended precision instruction is executed and the feature is not present in the system. This interrupt causes the processing program to call in a subroutine to handle extended precision operations. Extended precision divide is always simulated, because the extended precision facility does not include such an instruction.

Interval timer. The interval timer will be supported for the same functions as it is currently except for time of day values.

Time of day clock. This clock will be supported in MFT and MVT environments for time of day requests made by system and user tasks via the TIME macro. At IPL, the operator will have the option of validating the clock time and correcting an invalid value.

Byte boundary alignment. The programmer can byte-align binary and floating-point data in Assembler Language programs, in ANS COBOL programs (by omitting the SYNCHRONIZED clause), in PL/I programs (by specifying the UNALIGNED attribute), and in FORTRAN programs. However, OS still expects the parameters passed to it to be properly aligned and the high level language translators still align unaligned binary and floating-point data prior to its use.

7080, 7070/7074, 709/7090/7094/7094II Compatibility Features. Three separate, mutually exclusive emulator programs will be provided. Emulator programs are discussed in Section 40.

Channels. One or two 2870 Multiplexer Channels and up to twelve high-speed channels (2860 and 2880) in a system will be supported.

Both selector and block multiplexer mode of operation will be supported for 2880 channels. During IPL, block multiplexer mode is established via a control register bit setting (discussed in Section 10:20) if the user requested block multiplexer mode for at least one 2880 channel at system generation. The operator does not have the option of changing this mode at IPL, nor does the control program alter the mode during system operation.

The 3330 facility. The 3330 facility will be supported as an I/O device for most of the same functions as is the 2314 facility and by ASP and HASP. The error recovery routine provided for the 3330 will include support of the new hardware error correction features.

RPS will be supported in MFT and MVT environments as follows.

- The stand-alone seek issued within the I/O supervisor (IOS) will be eliminated for RPS devices (3330 facilities). IOS will continue to issue stand-alone seeks for direct access devices without RPS. IOS will also be capable of recognizing the channel available interrupt.

- Access method support of RPS commands (SET SECTOR and READ SECTOR) will be provided by:

  QSAM and BSAM for all record formats and functions except the undefined track overflow record format
  ISAM for all updating and verification and for QISAM sequential processing
  BPAM for processing directory and member processing operations where possible
  BDAM for direct retrieval and update of fixed-length standard and VBS format records without key and for write verification of all BDAM format records

- End-of-volume (EOV) routines will support concatenation of data sets residing on RPS and on non-RPS devices. The control program will ensure that an RPS access method is used for drives with the feature and that a non-RPS method is used for drives without it.

- Any system utility, data set utility, or IBM-supplied processing program (such as a language translator) that uses the sequential access methods will support RPS. In addition, IEHMOVE, IEBCOPY, and the initialize/dump/restore utility (IEHDASDR) will include RPS support for 3330 facilities.

- The Sort/Merge program product supports RPS for 3330 intermediate work devices.

- Where appropriate, RPS commands for access to SYSRES data sets will be supported by:

  Data set catalog routines
  Direct access device space management (DADSM) routines (for DSCB processing)
  STOW, BLDL, and FIND processing of program library (PDS) directory entries
  OPEN/CLOSE/EOV processing of JFCB's in the job queue
  Routines that access the job queue

  Note that RPS command support is not provided for:

  Access to TSO swap data sets
  Program fetch
  The telecommunications access method (TCAM) for message queues
  The stand-alone disk initialization and alternate track assignment routines (DASDI and IEHATLAS)

The 2305 facility. The 2305 facility will be supported as an I/O device for most of the same functions as is 2301 Drum Storage. Note that specification of exchange buffering for a QSAM data set on either a 2305 or 2330 facility results in a default to simple buffering.

Rotational position sensing for 2305 facilities will be supported as discussed for the 3330 facility except that an arm positioning seek is not required on 2305 modules and the 2305 is not supported as an intermediate work unit by Sort/Merge programs. In addition, the 2305 is not supported by ASP, HASP, or for TCAM message queues. Multiple requesting will be handled by the I/O supervisor, which will initiate up to seven channel programs on one 2305 module at a time. (The eighth logical device address is reserved for control purposes.) Multiple I/O operations to the same data set will be initiated if the data set is not being sequentially processed. (One operation must complete successfully before the next can be initiated during sequential processing.)

3211 Printer with tapeless carriage. The 3211 Printer, with or without the 18 additional print positions, will be supported by QSAM and BSAM for exactly the same functions as is the 1403 Printer and by ASP and HASP. The control program will handle loading of the forms control buffer (FCB) with carriage control images. This support is similar to that provided for Universal Character Buffer (UCB) loading.

The user can define one or more default FCB images at system generation time. Two IBM-supplied default images are included automatically. All other FCB images to be used must be defined by the user and placed in SYS1.SVCLIB, as is the case with UCB images. User-supplied default images must be identified as defaults. The FCB image to be used by a processing program can be specified in the 3211

Printer DD statement included for the job step and will be loaded into the FCB by the control program.

The FCB image currently loaded can also be changed by the programmer during execution of the processing program by use of an Assembler Language macro. If the DD statement does not specify an FCB image and the image currently loaded is not one of the defaults specified at system generation, the operator is requested to specify the FCB image to be used.

The FCB image (and the UCS character image) to be loaded for a 3211 Printer used by an output writer can be indicated in the output writer procedure or in the START output writer command issued by the operator. FCB and UCB images can also be specified in the SYSOUT DD statement for a data set that is to be printed by the output writer, and they will be loaded into the 3811 control unit prior to the printing of the data set.

Any time the FCB parameter is used, as described above, the user can specify that operator verification of forms alignment is to be requested by the control program via a console message when the buffer is loaded. The operator must respond to this message.

The 3211 error recovery routine will retry a print operation after a parity check occurs in the UCB or print line buffer if QSAM is used and if three I/O buffers are provided for the printer data set. When the operation is retried, the 3811 control unit insures that only the print positions that did not print correctly the first time are reprinted.

Operator console. The alphameric keyboard and CRT will be supported as the primary operating system console. The multiple console support (MCS) and device-independent display operator console (DIDOCS) options of OS must be included in the operating system to support this console.

The 3803/3420 Magnetic Tape Subsystem. This tape subsystem will be supported as an I/O device for the same functions as 2400-series tape units. This includes support of tape switching, Seven-Track, and Dual Density features. (Note that 200-BPI density tapes are not supported because the Seven-Track feature includes only 556 and 800-BPI densities.) The two-channel switch is supported for alternate paths switching. The RESERVE and RELEASE commands are not supported.

SECTION 40: EMULATOR PROGRAMS

40:05 FEATURES COMMON TO 7000-SERIES EMULATOR PROGRAMS

One of the significant new features available to Model 165 users is integrated 7000-series emulation under OS. Only stand-alone 7000-series emulators are available for the Model 65.

Three emulator programs, which use the three optional, mutually exclusive hardware compatibility features, are provided for the Model 165:

- 7070/7074 Emulator program - requires Compatibility Feature (#7117)
- 709/7090/7094/7094II Emulator program - requires Compatibility Feature (#7119)
- 7080 Emulator program - requires Compatibility Feature (#7118)

These emulator programs operate as processing programs in MFT and MVT environments and are alike in their basic functional design. Those features common to all three are discussed in this subsection.

The emulator programs supplied for the Model 165 are relocatable and thus can operate in one or more MFT partitions or MVT regions. Subject to the availability of system resources, any number of emulator jobs of the same type (7080, 7074, or 7094) can execute concurrently with System/370 jobs in the same Model 165. Emulator and System/370 jobs can be intermixed in the input stream, since emulator job scheduling, initiation, and resource allocation are handled by OS job management routines. I/O operations are handled by OS data management. Emulator jobs are executed by job priority, as is any OS job.

Integrated emulation provides a number of advantages over stand-alone emulation that can increase system throughput. The ability to execute 7000-series jobs under operating system control offers emulation users the benefits of multiprogramming and OS facilities. The advantages of integrated emulation are:

- Significantly better resource utilization, since System/370 native mode and 7000-series emulator jobs can be multiprogrammed. Stand-alone emulators normally use only a portion of the hardware resources available in the system.
- Standardized and simplified job accounting and job scheduling. The latter reduces the number of IPL's required, because switching from operating system to stand-alone emulation mode of system operation is not required.
- The ability to extend or add applications such as graphics, teleprocessing, time sharing, etc., because a dedicated emulation environment is no longer required and more system resources are available in a given time period.
- The ability to process emulated 7000-series job data sets, using both System/370 and emulated 7000-series programs. Existing tape files can be converted to a standard OS data format by using the IBM-supplied Tape Preprocessor formatting program.
- Device independence for emulator-supported I/O devices used by emulated 7000-series programs that handle data sets in OS VBS data

format. OS access methods, QSAM and BSAM, are used to handle I/O operations so that new functions and I/O device support added to the access method routines used by the emulator programs are automatically available to emulated 7000-series jobs. Tape and unit record 7000-series files can be emulated on System/370 direct access devices.

Each of the three emulator programs provided for the Model 165 uses simulation routines, one of the hardware compatibility features, the Model 165 instruction set, and OS supervisor and data management routines to emulate 7000-series program operations. Figure 40.05.1 shows the general layout of an emulator program partition or region.

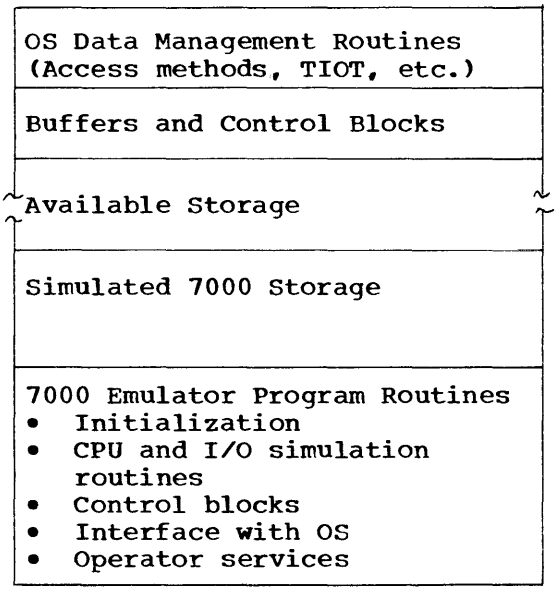

Figure 40.05.1. Partition or region layout for a 7000-series emulator program job step

The specific emulator program to be used by an installation must be constructed via an emulator generation procedure, which produces control statements required to link-edit the desired emulator modules and place the emulator program in SYS1.LINKLIB. Emulator program routines and two tape formatting programs (Tape Preprocessor and Tape Postprocessor) are distributed on a restore tape. The following must be done to include one of the emulators in an OS operating system for the Model 165:

- The fact that an emulator will be used with the operating system generated must be specified in the input required to generate the OS control program. This will cause the required emulator routines to be included in the generated operating system.

- The emulator restore tape must be obtained from PID and an emulator program with desired facilities must be generated. More than one version of a given 7000-series emulator program can exist in an operating system.

The microcode required by the compatibility feature ordered is distributed to the user on a magnetic disk cartridge as part of the system microcode that is loaded in WCS when the power-on button is pushed.

The emulator program generated can be used to execute existing 7000-series programs without change, subject to the following conditions:

1. Time-dependent programs being executed on a 7000-series system or emulated on a Model 65 may not execute properly. Provision has been made to allow some time-dependent programs to be emulated correctly (see the appropriate emulator reference manual for details).

2. Programs with undetected programming errors give unpredictable results.

3. If programs that use unsupported features or I/O devices (as described for each emulator in successive subsections and in emulator reference manuals) are to be emulated, they must be modified to conform to the support provided by the specific emulator program.

The Tape Preprocessor and Tape Postprocessor formatting programs supplied to Model 165 7000-series emulator users operate as processing programs and can be executed with any OS control program generated with the emulator macro specified. The Tape Preprocessor operates in a program area of 4K bytes plus I/O buffer requirements and accepts as input seven- and nine-track tape in 7000-series format. It produces as output spanned variable-length (OS VBS) format data that can be written on seven- or nine-track tape or on direct access storage. Input records longer than 32,755 bytes are reblocked, since OS BSAM cannot handle a physical data block longer than 32K bytes.

The Tape Postprocessor operates in a program area of 5K bytes plus I/O buffer requirements and performs the reverse of the Tape Preprocessor. The postprocessor program is useful when a copy of a data set in OS VBS and another in 7000 format are required. (The 7000-series emulator programs accept as input and produce as output both the formats handled by these two tape formatting programs.)

The tape formatting programs can handle 200, 556, 800, and 1600 BPI densities, mixed density tape volumes, and even, odd, and mixed parity tapes.

While existing tape files with blocks longer than 32K bytes must be preprocessed, conversion to VBS format offers the user the following:

- The ability to emulate tape data sets on direct access devices

- The ability to share among several concurrently executing emulator jobs read-only files maintained on direct access devices, such as 7000 system libraries.

- The ability to process VBS format data sets with both OS and emulated 7000-series programs. (Note that OS programs must be written with knowledge of the VBS format of preprocessed tapes.)

- The ability to increase emulator job performance by reblocking 7000-series format tape files with short blocks, assuming that enough processor storage is available

- The ability to reduce processor storage buffer requirements by reblocking files with very large blocks

The following subsections discuss the three emulator programs. They discuss features and I/O device support, Model 165 configuration requirements, job submission, conversion requirements, and performance.

40:10 7070/7074 EMULATOR PROGRAM

The OS 7070/7074 Emulator program for the Model 165 (hereafter referred to as the 7074 Emulator program) requires a Model 165 with the 7074 Compatibility Feature (#7117), 512K or more of processor storage, and enough I/O devices for the operating system and emulated devices. The emulator program requires a minimum processing partition or region of 188K in which to operate. This minimum supports a 7074 configuration of 10K words, two channels, seven tape drives per channel, and single buffering with 2000-byte buffers for each tape data set. (Note that buffering, I/O device configuration, performance, etc., will often necessitate use of an emulator partition or region larger than 188K.)

Table 40.10.1 lists 7074 system features that are supported and Table 40.10.2 indicates those that are unsupported. (See Emulating the 7070/7074 on the IBM System/370 Model 165 Using OS/360, GC27-6948, for an indication of how unsupported instructions are handled by the emulator.) The 7074 emulator also can emulate 7070 programs that do not contain elements that are incompatible with the 7074, as discussed in IBM 7070/7074 Principles of Operation (GA22-7003). The Model 165 7074 emulator supports the same facilities as the Model 65 stand-alone 7074 emulator except that 7074 unit record devices (7500 Card Reader, 7550 Card Punch, and 7400 Printer) are supported by the Model 65 emulator but not by the Model 165 emulator.

The 7074 Emulator program accepts and produces two tape data formats using BSAM:

1. 7074 format tapes, written in BCD, that are written by a 7074 system, a 1401 system (or a 1401 emulator), the stand-alone 7074 emulator, the Model 165 integrated emulator, or the Tape Postprocessor program. These tapes can contain segment marks and embedded tapemarks but cannot contain binary (odd parity) format data. Mixed density, seven-track input tapes can be handled if single buffering is used and the tape density specified in the DD statement is that of the first record on the tape.

   The user is responsible for end-of-volume switching for multivolume 7074 format tape files. Tapes in 7074 format can be processed on nine-track and seven-track 2400 tape drives. The latter must be attached to a control unit with the Seven-Track Compatibility Feature. Tapes in 7074 format must also be designated as unlabeled as far as OS is concerned and their record format must be specified as undefined.

2. Spanned variable-length (OS VBS) format tapes, written in EBCDIC, that are produced by the IBM Tape Preprocessor program or as a result of 7074 emulation on a Model 165. These tapes can be read by both System/370 and emulated 7074 programs. End-of-volume switching is handled by OS data management and is therefore transparent to the 7074 programmer. (Thus, if a 7074 program depends on an end-of-volume indication, 7074 format should be used.)

VBS format tape volumes can be processed on a nine-track 2400 tape unit or a 2400 tape unit with a seven-track head and the Data Converter feature. These volumes can have unlabeled data sets or standard OS labels in additon to 7074 labels. Note that VBS format files also can be processed on direct access devices.

Table 40.10.1. 7074 hardware and I/O devices supported by the Model 165 7074 Emulator program

7074 Features	Model 165 Equivalent
•7074 System (or 7070) with 10K words, up to four channels, and floating-point arithmetic	•Model 165 with the 7074 Compatibility Feature, 512K or more, channels (2870, 2860, or 2880), and an MFT or MVT control program that includes a 7074 Emulator program and the interval timing, ATTACH, and IDENTIFY options
•7150 Console All features emulated except: LOG operating key INQUIRY ONLY status key PRIORITY CONTROL dials for unit record	•OS system console
•7501 Console Card Reader	•1442 Card Read Punch •2501 Card Reader •2520 Card Read Punch •2540 Card Read Punch •Any I/O device supported by OS QSAM
•729 II/IV/V/VI Magnetic Tape Units Binary (odd parity) data and formats that depend on use of the TRA instruction are not supported.	•Any System/370 tape unit or direct access device supported by OS BSAM Data must be in VBS format in order to be processed on a direct access device.
•Data check facilities: Accumlator overflow halt mode Field overflow Sign change Exponent overflow Program check Arithmetic checks Validity checks for valid operations codes, valid priority control masks, and valid floating-point data format	•Emulator program

Table 40.10.2. 7074 I/O devices and features not supported by the 7074 Emulator program for the Model 165

7907 Data Channel and I/O devices it controls
7900 Inquiry Station
7631 File Control and 1301/2302 Disk Storage
7500 Card Reader
7550 Card Punch
7400 Printer
7603 Input/Output Synchronizer
Tape read all alpha (TRA) instruction
Read binary tape instruction
Tape read or write from 7074 locations 9990-9999
729 Tape Switching feature
Unit record priority interrupts
Unit record signal
Additional storage above 10K words (and associated instructions)
Interval timer
Customer engineer diagnostic instructions
Optional features other than floating-point arithmetic

OS job control and 7074 emulator control statements must be present for each emulated 7074 object program. Existing 7074 programs and their control statements need not be modified unless they use facilities not supported by the Model 165 7074 Emulator program.

The 7074 object programs to be emulated can be placed in the input stream, on tape, or in a sequential data set on disk (including a partitioned data set). A DD statement in the input stream describes their location. The required 7074 emulator control statements can be provided via the operator console or in a card, tape, or disk sequential data set identified by a specially named DD statement (SYSEMCTL). They can also be placed in the input stream or in a partitioned data set (as a member).

The 7074 Emulator program operator commands provided can be entered via the operator's console or the input stream (located in the emulator control statement data set). These commands provide simulation of 7074 commands on the System/370 operator's console. In addition, they allow the operator to request a listing of an entire emulator program area or of emulated 7074 main storage within an emulator partition or region only.

If multiple console support (MCS) is included in the OS control program generated, emulator program messages can be routed to a specific console device so that emulation messages are isolated.

Assuming that the 7074 (or 7070) programs to be emulated do not use facilities or I/O devices unsupported by the Model 165 7074 Emulator program, the following is required:

- OS job control and any necessary Model 165 7074 emulator control statements must be supplied for each 7074 object program. Changes to existing object programs and 7074 control statements are not required. (This is true for 7074 programs being executed on a 7074 system or emulated on a Model 65.)

- Tape volumes containing files with records longer than 32,755 bytes must be preprocessed by the tape formatting program. Tape written in either seven- or nine-track 7074 mode can be used without modification. (Consideration should be given to reblocking files with short blocks in order to improve execution time and reblocking files with long blocks to reduce processor storage buffer requirements.)

- Those 7074 programs that use unsupported I/O devices or facilities can be modified to remove the unsupported item and then emulated, or these programs can be rewritten to run in System/370 mode.

The internal performance (that is, the speed of performing 7074 CPU instructions weighted by frequency of use) of the 7074 Emulator program for the Model 165 is approximately three times that of a 7074 system. The throughput achieved by the integrated 7074 emulator operating on the Model 165 versus that obtained using an 7074 is dependent on the characteristics of the 7074 program being emulated, the hardware resources available to the 7074 emulator, and the number of other jobs operating concurrently with the emulated 7074 job(s). Emulator performance is improved when each 7074 channel is emulated with a single Model 165 channel (2860, 2880, or 2870 selector subchannel). Total system throughput should be improved by use of integrated, rather than stand-alone, emulation.

40:15 709/7090/7094/7094II EMULATOR PROGRAM

This emulator is referred to in this subsection as the 7094 Emulator program. The use of 7094 in the text refers to all 709X systems emulated unless otherwise indicated.

A Model 165 with the 7094 Compatibility Feature (#7119), 512K (MFT) or 1024K (MVT) or more of processor storage, and enough I/O devices for the operating system and emulated 7094 programs is required. An online printer dedicated to the 7094 emulator is recommended. The 7094 Emulator program requires a minimum partition or region size of 374K in which to execute. This minimum supports two channels with ten tape units each and one 1560-byte buffer per tape data set. (Note that buffering, I/O device configuration, performance, etc., will often necessitate use of an emulator partition or region larger than 374K.)

Table 40.15.1 lists the 7094 system features that are supported and Table 40.15.2 indicates those that are unsupported. The facilities supported by the Model 165 7094 Emulator program are the same as those provided by the integrated 7094 Emulator program for the Model 85. The hardware and I/O devices supported by the Model 165 7094 emulator are the same as those supported by the Model 65 stand-alone 7094 emulator with two exceptions. The Model 65 stand-alone 7094 emulator supports 704 emulation and mixed density tape, but the Model 165 7094 emulator does not. (The Model 85 emulator does not support 704 emulation or mixed density tape either.)

The 7094 Emulator program accepts and produces two data formats using BSAM:

1. 7094 format tapes written by the Tape Postprocessor program, a 7094 system, a 1401 system or 1401 emulator, or the 7094 emulator for the Model 65, 85, or 165. These tapes can contain embedded tapemarks and mixed mode files, that is, BCD (even parity) and binary (odd parity) data.

2. Spanned variable-length (OS VBS) format tapes that are produced by the Tape Preprocessor program or one of the IBM emulators. Embedded tapemarks and mixed parity data cannot be handled in this format, but BCD (even parity) characters are represented by using equivalent EBCDIC characters, and a parallel set of characters is used to represent binary data (odd parity). Special blocks are created to simulate tapemarks.

VBS format tape volumes can contain OS standard labels in addition to 7094 labels and can be processed by both OS and emulated 7094 programs. A nine-track 2400-series tape unit or a seven-track tape

unit with the Data Converter feature is required to process VBS format tape volumes. Note that VBS format files can also be processed on direct access devices.

Table 40.15.1. 7094 hardware and I/O devices supported by the 7094 Emulator program for the Model 165

7094 Feature	Model 165 Equivalent
•709/7090/7094/7094II system with 32K words of main storage and up to eight channels (A through H)	•Model 165 with the 7094 Compatibility Feature, 512K (MFT) or 1024K (MVT) or more, channels (2860, 2880, or 2870), and a MFT or MVT operating system that includes the interval timing, ATTACH, and IDENTIFY options and a 7094 Emulator program
•729 Magnetic Tape Units (Mixed parity files - with even BCD and odd binary parity data - are supported but mixed density files are not.)	•Any System/370 tape unit supported by OS BSAM •Any System/370 direct access device supported by OS BSAM if VBS format is used. (Only single-volume disk data sets are supported.)
•711 Card reader	•Any card reader or SYSIN device supported by OS QSAM. The Card Image feature is required to handle 7094 binary card input. Binary card input cannot be handled by the regular reader interpreter. MVT users can invoke the ASB reader to read binary card input in the input stream. •Any System/370 direct access device supported by OS BSAM. (Only single-volume disk data sets are supported.)
•721 Card Punch	•Any card punch or SYSOUT device supported by OS QSAM. The Card Image feature is required on a card punch. •Any System/370 direct access device supported by BSAM. (Only single-volume disk data sets are supported.)
•716 Printer (Simulation is based on the assumption of standard SHARE 2 control board wiring.)	•Any printer or SYSOUT device supported by OS QSAM. •Any System/370 direct access device supported by OS BSAM. (Only single-volume disk data sets are supported.) •Any System/370 console. (Spacing and carriage control is not provided if the console is used.)

Table 40.15.2. 7094 I/O devices and features not supported by the 7094 Emulator program for the Model 165

7909 Data Channel and attached devices:
1414 Input/Output Synchronizer Model 6
7631 File Control and 1301, 1302 Disk Storage
7640 Hypertape Control and 7340 Hypertape Drives
7320 Drum Storage
7740 Communication Control System
7750 Programmed Transmission Control
Data Channel Switch
Direct data connection
Direct couple
External signal trap
740 CRT Recorder (709 device)
780 CRT Recorder (709 device)
7094 special and custom features
704 Compatibility Feature
SET DENSITY instruction (mixed density tapes)

System/370 and 7094 emulator jobs can be intermixed in the OS input stream, with special processing necessary sometimes for binary format jobs and data. OS job control and any required 7094 emulator control statements must be supplied for each 7094 program to be emulated. Existing 7094 object programs and control statements need not be modified unless they use facilities not supported by the Model 165 7094 Emulator program.

BCD format 7094 jobs and data can be placed in the input stream or in a card or tape data set. Binary format 7094 jobs in card form (object programs and data) can be placed in the input stream (in a card reader with the column binary feature) in an MVT environment if the automatic SYSIN batching (ASB) reader interpreter is used. Binary card input can also be read from a dedicated card reader with the column binary feature. Required emulator control statements can be entered via the operator's console or a data set identified by a specially named DD statement (SYSEMCTL). This data set can be in the input stream or in a card, tape, or disk sequential data set (including a member of a partitioned data set).

The 7094 Emulator program operator commands available can be entered via the operator's console or the input stream (located in the emulator control statement data set). These commands provide simulation of 7094 console operations and allow the operator to request dumps of emulator processor storage areas and to obtain certain displays (sense switch settings, register values, etc.). If multiple console support (MCS) is included in the OS control program generated, emulator program messages can be routed to a specific console device (assigned routing code 12) so that emulation messages are isolated.

Assuming that the 7094 programs to be emulated do not use facilities or I/O devices unsupported by the Model 165 7094 emulator, the following is required:

- OS job control and any required Model 165 7094 emulator control statements must be supplied for each 7094 object program to be emulated. Changes to BCD format data is not required. If binary card input is not handled by card reader operations, a utility program must be written to prepare an acceptable tape from the binary input. A user-written data set writer is required to print and punch 7094 files written to SYSOUT data sets, if any. (See Emulating the 709, 7090, 7094, 7094II on the IBM System/370 Model 165 Using OS/360, GC27-6951, for suggestions regarding these utilities.)

- Tapes containing files with records longer than 32,755 bytes (5461 7094 words) must be preprocessed by the tape formatting program. Tapes written in seven- or nine-track 7094 mode can be used without modification. (Consideration should be given to reblocking files with short blocks in order to improve execution performance and reblocking files with long blocks to reduce processor storage buffer requirements.)

- Those 7094 programs that use unsupported I/O devices or facilities can be modified to remove the unsupported facility and then emulated, or they can be rewritten as System/370 programs.

The internal performance (that is, the speed of performing 7094 CPU instructions weighted by frequency of use) of the Model 165 7094 Emulator program is approximately 1.5 times that of the IBM 7094II. The throughput of emulated 7094 jobs executed on the Model 165 will vary depending on job characteristics, resources available, etc. Total system throughput increases should be achieved by operating 7094 emulator jobs in a multiprogramming environment.

40:20 7080 EMULATOR PROGRAM

The 7080 Emulator program for the Model 165 requires a Model 165 with the 7080 Compatibility Feature (#7118), 512K or more of processor storage, and enough I/O devices for the operating system and emulated devices. In an MVT environment a 1024K minimum system is required to emulate a 7080 system with 160,000 positions of main storage. The emulator program requires a minimum processing partition or region of 232K for emulation of a 7080 system with 80K. A 364K partition or region is required to emulate a 7080 system with 160K. These two minimums support two channels and 14 tape units with one 2000-byte buffer assigned to each tape data set. (Note that other considerations, such as double buffering, I/O device configuration, performance, etc., will often necessitate use of an emulator partition or region larger than these minimums.)

Table 40.20.1 lists 7080 system features that are supported and Table 40.20.2 indicates those that are unsupported. The 7080 emulator also can execute 705 I, II, and III programs that can be executed on a 7080, subject to the considerations and limitations outlined in the IBM 7080 Principles of Operation (GA22-6560) and in this document. The Model 165 emulator supports the same facilities as the Model 65 stand-alone 7080 emulator.

The 7080 Emulator program accepts and produces two tape data formats, using BSAM:

1. 7080 format tapes, written in BCD, that are written by a 7080 system, the Model 65 stand-alone 7080 emulator, the Model 165 integrated 7080 emulator, or the Tape Postprocessor program. These tapes can contain imbedded tapemarks but cannot contain the reader storage mark character. Mixed density, seven-track input tapes can be handled if single buffering is used and the tape density specified in the DD statement is that of the first record on the tape.

   Tapes in 7080 format can be processed on nine-track tape drives and 2400 tape drives with a seven-track head attached to a control unit with the Seven-Track Compatibility Feature. They must also be designated as unlabeled as far as OS is concerned and their record format must be specified as undefined.

2. Spanned variable-length (OS VBS) format tapes, written in EBCDIC, that are produced by the Tape Preprocessor program or as a result of 7080 emulation on a Model 165. These tapes can be read by both System/370 and emulated 7080 programs.

   VBS format tape volumes can be processed on a nine-track 2400 tape unit or a 2400 tape unit with a seven-track head and the Data Converter feature. These volumes can have unlabeled data sets or standard OS labels in addition to 7080 labels. Note that VBS format files also can be processed on direct access devices.

Table 40.20.1. 7080 hardware and I/O devices supported by the 7080 Emulator program for the Model 165

7080 Feature	Model 165 Equivalent
7080 system with 80,000 or 160,000 positions of storage, and up to four 7080 communication channels. (705 I, II, and III programs that can be executed on a 7080 system can be emulated also, subject to Model 165 7080 emulator restrictions.)	Model 165 with the 7080 Compatibility Feature, 512K or more, channels (2870, 2860, or 2880), and an MFT or MVT control program that includes a 7080 Emulator program and the interval timing, ATTACH, and IDENTIFY options. A 1024K system is required for emulation of a 7080 with 160,000 positions of main storage, using MVT.
7153 console	OS system console
7502 Console Card Reader	1442 Card Read Punch 2501 Card Reader 2520 Card Read Punch 2540 Card Read Punch Any I/O device supported by OS QSAM
729 II/IV/V/VI Magnetic Tape Units	Any System/370 tape unit or direct access device supported by OS BSAM. VBS format is required for direct access device emulation. 2400 tape units with the Seven-Track Compatibility Feature are required to handle 7080 mode tapes.

Table 40.20.2. 7080 I/O devices and features not supported by the 7080 Emulator program for the Model 165

7622 Signal Control, associated units, and related instructions
7908 Data Channel, associated devices, and features (including control storage bank 4) and related instructions (ENABLE COMPARE BACKWARD)
7080 RPQ features and related instructions
7340 Hypertape Drive
1301 and 2302 Disk Storage
ASU zoning contained in multiply, divide, and lengthen instructions
Customer engineering diagnostic instructions
Automatic restart indicator

OS job control and 7080 emulator control statements must be present for each emulated 7080 object program. Existing 7080 programs and their control statements need not be modified unless they use facilities not supported by the Model 165 7080 Emulator program.

The 7080 object programs to be emulated can be placed in the input stream, on tape, or in a sequential data set on disk (including a partitioned data set). A DD statement in the input stream describes their location. The required 7080 emulator control statements can be provided via the operator console or in a card, tape, or disk sequential data set identified by a specially named DD statement (SYSEMCTL). They can also be placed in the input stream or in a partitioned data set (as members).

The 7080 Emulator program operator commands provided can be entered via the operator's console or the input stream (located in the emulator control statement data set). These commands provide simulation of 7080 commands on the Model 165 operator's console. In addition, they allow the operator to request a listing of an entire emulator program area or of emulated 7080 main storage within an emulator partition or region only.

If multiple console support (MCS) is included in the OS control program generated, emulator program messages can be routed to a specific console device so that emulation messages are isolated.

Assuming that the 7080 (or 705 I, II, and III) programs to be emulated do not use facilities or I/O devices unsupported by the Model 165 7080 Emulator program, the following is required:

- OS job control and any necessary Model 165 7080 emulator control statements must be supplied for each 7080 object program. Changes to existing object programs and 7080 control statements are not required. (This is true for 7080 programs being executed on a 7080 system or emulated on a Model 65.)

- Tape volumes containing files with records longer than 32,755 bytes must be preprocessed by the tape formatting program. Tape written in either seven- or nine-track 7080 mode can be used without modification. (Consideration should be given to reblocking files with short blocks in order to improve execution time and reblocking files with long blocks to reduce processor storage buffer requirements.)

- Those 7080 programs that use unsupported I/O devices or facilities can be modified to remove the unsupported item and then emulated, or these programs can be rewritten to run in System/370 mode.

The internal performance (that is, the speed of performing 7080 CPU instructions weighted by frequency of use) of the 7080 Emulator program for the Model 165 is approximately 2 times that of a 7080 system. The throughput achieved by the integrated 7080 emulator operating on the Model 165 versus that obtained using an IBM 7080 is dependent on the characteristics of the 7080 program being emulated, the hardware resources available to the 7080 emulator, and the number of other jobs operating concurrently with the emulated 7080 job(s). Emulator performance is improved when each 7080 channel is emulated with a single Model 165 channel (2860, 2880, or 2870 selector subchannel). Total system throughput should be improved by use of integrated, rather than stand-alone, emulation.

SECTION 50: RELIABILITY, AVAILABILITY, AND SERVICEABILITY (RAS) FEATURES

50:05 INTRODUCTION

With the growth of more and more online data processing activities, as distinguished from traditional batch accounting functions, the availability of the data processing system becomes a very essential factor in company operations, and complete system failure is extremely disruptive. Because of the growing frequency of online processing and the fact that the System/370 Model 165 is designed to operate in such an environment, IBM has provided an extensive group of advanced reliability, availability, and serviceability features for the Model 165. These RAS features are designed to improve the reliability of hardware, to increase the availability of the computing system, and to improve the serviceability of system hardware components.

The RAS features of the System/370 Model 165 are designed to reduce the frequency and impact of system interruptions that are caused by hardware failure and necessitate a re-IPL as follows.

- More reliable components, such as integrated circuits with fewer connections, will be used to improve hardware reliability.
- Recovery facilities, both hardware and programming systems, not available on System/360 Models 65 and 75, are provided to reduce considerably the number of failures that cause a complete system termination. This permits deferred maintenance.
- Repair procedures include more online diagnosis and repair of malfunctions concurrently with normal job execution in a multiprogramming environment in order to reduce the effect of such repairs on system unavailable time.

Each RAS feature, recovery or repair, is discussed in the remainder of this section.

The following recovery features are implemented in hardware:

- CPU retry of most failing CPU operations, including those caused by a buffer malfunction
- ECC validity checking on processor storage to correct all single-bit errors
- I/O operation retry facilities, including channel retry data and channel/control unit command retry procedures, to correct failing I/O operations
- Expanded machine check interrupt facilities to facilitate better error recording and recovery procedures

The following recovery features are provided by programming systems:

- Recovery management support (RMS) to handle the expanded machine check interrupt and channel retry data. Model 165 MCH and CCH routines are provided for OS MFT and MVT only.
- Error recovery procedures (ERP) to retry failing I/O device and channel operations

- OBR and SDR routines to record statistics for all temporary and permanent I/O errors
- Environment recording, edit, and print program (EREP) to format and print error log records
- I/O RMS routines, alternate path retry (APR), and dynamic device reconfiguration (DDR) to provide additional recovery procedures after channel or I/O device failures
- Advanced checkpoint/restart and warm start facilities to simplify and speed up system restart procedures after a failure necessitates a re-IPL

The following repair features are provided:

- Online Test Executive Program (OLTEP) and Online Tests (OLT's) that execute under OS control and provide online diagnosis of I/O device errors for most devices that attach to the Model 165
- Processor Logout Analysis program that operates under OS control to analyze machine check error records in order to determine suspected malfunctioning field-replaceable units
- System Test, Channel Test, CPU Test, and Storage Test stand-alone diagnostic programs to identify failing hardware units
- Microdiagnostics for customer engineer use to locate the field-replaceable unit within a malfunctioning component

These aids are designed to enhance system availability. In many cases, the system can run in a degraded mode so that maintenance can be deferred to scheduled maintenance periods. When solid failures do occur, their impact can be reduced by faster isolation and repair of the malfunction than is possible currently.

50:10 RECOVERY FEATURES

Additional hardware that attempts correction of most hardware errors without programming assistance has been included as part of the basic Model 165 system. The control program can be notified, via an interrupt, of both intermittent and solid hardware errors so that error recording and recovery procedures can take place.

AUTOMATIC CPU RETRY

Detected CPU hardware errors, except those that occur during execution of certain instructions that have passed beyond a threshold point, can be retried automatically by CPU retry hardware. A mask bit in a control register determines whether the CPU retry function is enabled or disabled. If enabled, retry occurs after instruction errors, after failures that occur during interrupt time when status information is being saved, after errors that occur during status saving for I/O instructions, etc. An I/O instruction, such as START I/O or TEST I/O, can be retried automatically by the hardware without an intervening I/O interrupt if the instruction has not proceeded beyond an established threshold point.

CPU retry also occurs when an instruction error results from a buffer malfunction, if the instruction is a retryable type. The buffer is bypassed while the instruction is retried so that processor storage

is referenced directly. If the retry is successful, operations continue as usual. An interrupt occurs to indicate a successful CPU retry. The fact that a buffer failure occurred is indicated as well.

CPU retry is accomplished by additional microprogram routines and hardware included in the Model 165. The failing CPU operation is retried by the microprogram up to seven times before it is determined that the error is uncorrectable. Most instructions must be completely reexecuted. Therefore, all data required for a retry is saved by the execution microprogram. However, certain SS-format instructions will be retried from the point of successful execution. In this case, the execution microprogram saves the status data necessary to restart at the proper byte.

When enabled, a machine check interrupt takes place after a CPU error occurs and is retried. If CPU retry was successful, and no buffer error occurred, the failure need only be recorded. If a buffer failure caused the error or if the retry was unsuccessful, programmed recovery procedures are required in addition to error recording.

The CPU retry feature provides the Model 165 with the ability to recover from intermittent CPU failures that would otherwise cause a system halt and necessitate a re-IPL or that would cause an executing program to be terminated. Corrected errors are logged (by MCH) for later diagnosis during scheduled maintenance periods. System availability is thereby increased.

Retry of failing CPU operations on Models 65 and 75 is not provided by system hardware. Instruction retry after a machine check is provided for the Model 65 for some instructions by the machine check handler (MCH) routine. (An MCH routine is not provided for the Model 75.)

ECC VALIDITY CHECKING ON PROCESSOR STORAGE

The ECC method of validity checking on processor storage provides automatic single-bit error detection and correction. It also detects all double-bit and most multiple-bit processor storage errors but does not correct them. Checking is handled on an eight-byte basis, using an eight-bit modified Hamming code, rather than on a single-byte basis using a single parity bit. However, parity checking is still used to verify other data in a Model 165 system that is not contained in processor storage. Models 65 and 75 use parity checking for main storage data verification.

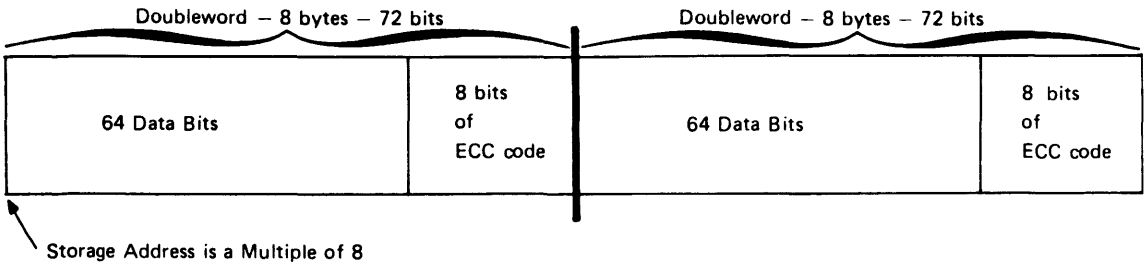

Figure 50.10.1. Data representation used in Model 165 processor storage

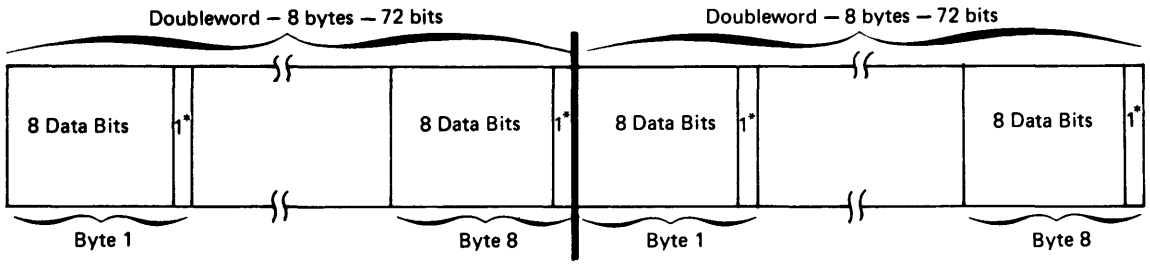

*Parity Bit

Figure 50.10.2. Data representation used in Models 65 and 75 and in the Model 165 in other than processor storage

Data enters and leaves processor storage through the storage adapter unit, which performs ECC validity checking on each doubleword. When a doubleword (72 bits) is fetched from processor storage, the storage adapter unit checks the 8-bit ECC code to validate the 64 data bits. If the data is correct, the adapter unit generates the appropriate parity bit for each of the 8 data bytes and reformats the doubleword to look as shown in Figure 50.10.2. If a single-bit error is detected, the identified data bit in error is corrected automatically by the corrector unit in the storage adapter. The corrected doubleword is sent back to processor storage and on to the storage control unit. When a doubleword is to be placed in processor storage, the storage adapter unit strips the 8 parity bits, constructs the necessary 8-bit ECC code, and appends the code to the 64 data bits. The 72 bits are then stored as shown in Figure 50.10.1.

When a single-bit storage error is detected and corrected during the execution of an instruction or I/O operation, a soft machine check latch is turned on and execution continues. At the completion of the operation a machine check interrupt occurs to allow error recording to be done unless ECC correction interrupts have been disabled.

When a double- or a multiple-bit processor storage error involving the CPU is detected, a machine check interrupt occurs. The MCH routine logs the error and attempts recovery procedures. When a double- or multiple-bit processor storage error occurs during an I/O operation, both a machine check and an I/O interrupt occur so that both error recording and I/O retry procedures can be executed.

An ECC mode bit controls whether full recording or quiet mode is in effect for corrected single-bit processor storage errors. When quiet mode is in effect, a machine check interrupt does not occur after the successful correction of a single-bit error.

The ECC feature increases Model 165 system availability by permitting system operation to continue normally after single-bit storage errors occur and are corrected. Any main storage errors on a Model 65 will necessitate the termination of system operation or of the program involved if the damaged program module cannot be refreshed. (Both control program and user modules can be refreshed by Model 65 MCH. Only SER routines are provided for the Model 75.)

I/O OPERATION RETRY

Channel retry, command retry, and CPU retry features are provided to reduce the number of abnormal program terminations and unscheduled system halts that occur because of channel errors.

- Channel Retry

  This feature has been implemented to insure that most failing channel operations can be retried by error-handling routines. When a channel error (channel control or interface control check) or a CPU error associated with a channel operation occurs, both the channel status word (CSW) and channel logout data are stored during the I/O interrupt.

  Channel logout data for the Model 165 consists of 28 words. It is stored beginning at the location pointed to in word 172 in the fixed area if the I/O extended log mask is enabled. The channel logout data provides additional, more exacting status information about the channel failure. This data is formatted by the CCH routine and passed to a device-dependent error recovery routine to be used in the retry of the failing I/O operation.

- Command Retry

  Command retry is a channel/control unit procedure that can cause an improperly executed command in a channel program to be retried automatically by hardware so that an I/O interrupt and programmed error recovery are not required. An indication is presented when the control unit recognizes this situation. The 2870 Multiplexer and 2860 Selector channels will not perform a command retry.

  The command retry feature is implemented in the control unit of the 3330 and 2305 facilities and was discussed in Section 20.

- CPU Retry

  If an error occurs during the execution of an I/O instruction, such as START I/O, TEST I/O, TEST CHANNEL, etc., the hardware determines whether or not the automatic CPU retry threshold for the particular instruction has been passed. If instruction execution has not gone beyond predetermined points, then the instruction is retried automatically by the CPU retry hardware without programming assistance. A machine check interrupt is taken for recording purposes at the completion of a successful retry. If the instruction cannot be retried because it has passed beyond the retry threshold point, an I/O interrupt is taken and the appropriate device-dependent error recovery routine is scheduled to take the required recovery action. For example, if an error in the execution of a START I/O instruction occurs before the I/O device becomes involved, a CPU retry is possible.

EXPANDED MACHINE CHECK INTERRUPT FACILITIES

Implementation of the machine check level of interrupt on the System/370 Model 165 has been expanded in order to enhance error recording and error recovery procedures. Programming support of the extended machine check interrupt is provided by the Model 165 MCH routine.

The machine check interrupt facilities of the Model 165 differ from those of Models 65 and 75 as follows.

- Four subclasses of machine check interrupt are defined.
- A machine check interrupt occurs to permit the recording of errors corrected by the hardware (soft machine check interrupt) as well as to allow recovery routines to handle errors that cannot be corrected by hardware (hard machine check interrupt).
- Machine check interrupt masking is expanded to handle selective disabling and enabling of the interrupt subclasses defined, the CPU retry facility, and the occurrence of extended logouts.
- The size of the fixed storage area in lower processor storage is increased to accommodate the storing of additional machine status and diagnostic information when a machine check interrupt occurs.
- A hard stop error condition is defined that causes the Model 165 system to stop functioning immediately because the nature of the machine malfunction prevents valid processing from continuing.
- A hang detect check is implemented that causes a machine check interrupt if a CPU operation has not been completed within 16.6 ms.

The Model 165 presents one of four subclasses of machine check interrupt, depending on the specific machine malfunction. Each interrupt subclass is maskable and causes either a soft machine check or a hard machine check interrupt when enabled. A soft machine check occurs after the hardware has been successful in correcting an error or when the error does not affect the interrupted program. This is done so that the failure can be recorded. System operation continues after the error is logged. For example, if an error occurs during the execution of an instruction and if the CPU retry hardware corrects the error by reexecuting the failing instruction, a soft machine check interrupt occurs at the completion of the successful execution.

A hard machine check occurs when hardware retry fails or is not possible. For example, if the CPU retry hardware has not corrected a failing instruction after seven retries, a hard rather than a soft machine check interrupt occurs after the last unsuccessful retry.

Figure 50.10.3 shows the layout of fixed processor storage in the Model 165. Fixed storage consists of three areas: the fixed locations in decimal addresses 0-127, the fixed logout area in locations 176-511, and the CPU extended logout area of 1000 bytes, which normally begins at location 512 unless the logout pointer is altered by programming.

Fixed locations 0-127 are identical in layout and content to these locations in System/360 models, with the exception of the EBCDIC/ASCII bit in the current PSW, which must be set to zero.

A logout to the fixed logout area (176-511) occurs when any type of machine check interrupt is taken. The data stored varies slightly

among System/370 models and is processed by recovery management routines. The fixed logout area data indicates the reason for the interrupt in the machine check code. The save areas in the logout area preserve the status of the system at the time of the machine check interrupt and contain the contents of the general purpose, the floating-point, and the control registers.

The model-dependent CPU extended logout area begins at the address specified in control register 15, which is set to decimal location 512 by an IPL or on a system reset. The length of this extended logout area on the Model 165 is approximately 1000 bytes. The occurrence of a logout to the CPU extended logout area is determined by the setting of two extended logout mask bits in control register 14. An extended logout can occur after a CPU error and when a hard machine check interrupt takes place. This data can be recorded by recovery management routines and the Processor Logout Analysis program can be used to process this information at a later time.

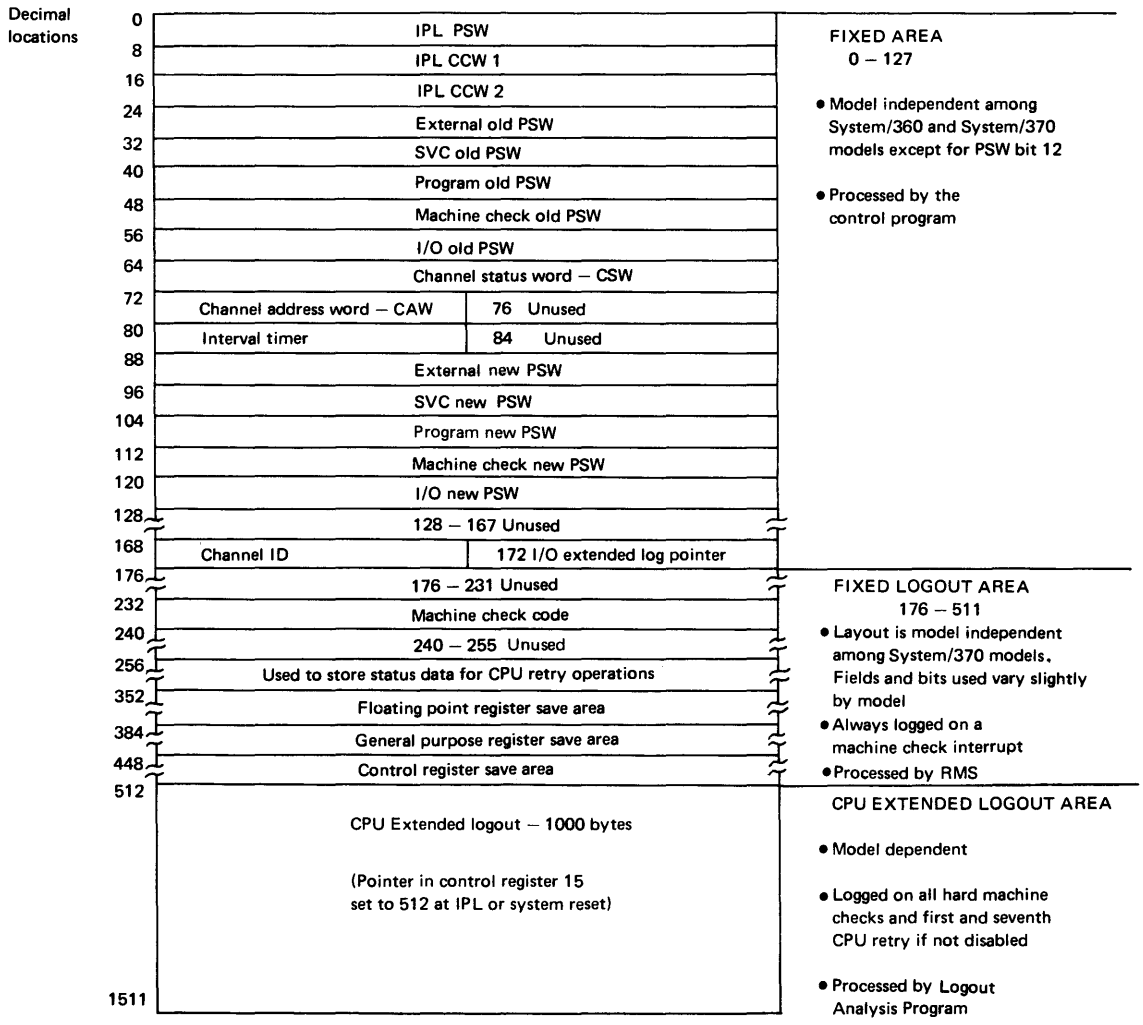

Figure 50.10.3. Model 165 fixed storage locations

Figure 50.10.4 illustrates the layout and the contents of the eight-byte machine check code stored in processor storage locations 232-239. The machine check code indicates which type of interrupt occurred, the validity of certain fields stored in the fixed logout area, and the length of the stored CPU extended logout area.

Table 50.10.1 lists the machine check subclasses defined for the Model 165. They are described in the discussion that follows.

The mask bits used to enable or disable each subclass for interrupts are indicated and the setting of the machine check code is discussed. PSW bit 13 and two other mask bits are used to enable and disable machine check interrupts. The recovery mask (R) and external mask (E) bits are contained in control register 14 and operate subject to PSW bit 13. If PSW bit 13 is disabled, then all machine checks are masked. If PSW bit 13 is enabled, then the settings of the two additional mask bits determine whether or not interrupts, other than System Damage, will be taken. Refer to Figure 50.10.4.

Table 50.10.1. Model 165 machine check interrupts

Subclass	Mask Bit(s)	Cause	Machine Check
System Damage	PSW 13	• Unretryable CPU error • Uncorrectable CPU error • Multiple-bit processor storage error • Storage protect key failure	Hard
System Recovery	PSW 13 and R	• CPU error corrected by retry • Single-bit processor storage error corrected by ECC	Soft
Time of Day Clock Damage	PSW 13 and E	• Error in time of day clock	Soft
External Damage	PSW 13 and E	• Error that did not affect the CPU, such as multiple-bit processor storage error during an I/O operation	Soft

Soft machine check interrupts on System/370 Model 165 are as follows:

- System Recovery - This interrupt occurs if both PSW bit 13 and the recovery mask bit are on. It is caused by a successful CPU retry or single-bit processor storage error correction.

  The SR bit in the stored machine check code (bit 2) is used to indicate the occurrence of an ECC single-bit error correction or a successful CPU retry of a failing CPU operation (including one caused by a buffer malfunction). The SC bit (bit 17) will be on as well if an ECC recovery occurred.

  Error recording and, possibly, buffer deletion are required after a System Recovery interrupt.

Fixed Logout Area Locations 232 - 239

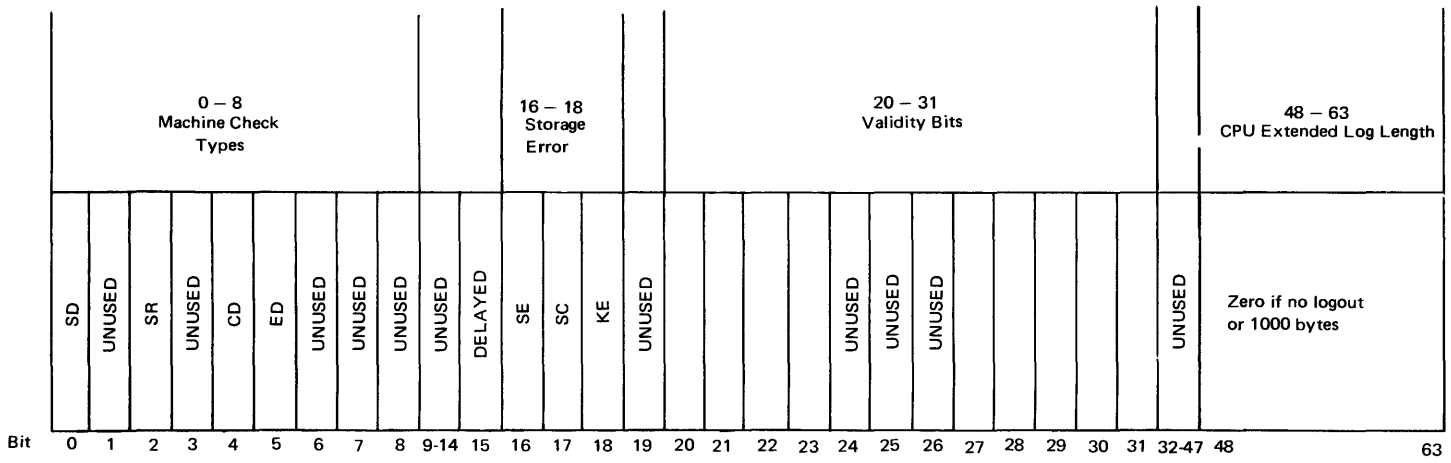

Bit	Inrerrupt Type
0	SD – System Damage
2	SR – System Recovery
4	CD – Time of Day Clock Damage
5	ED – External Damage

Bit	Error
15	Delayed Interrupt
16	Storage Error Uncorrected
17	Storage Error Corrected
18	Protection Key Error

Bit	Valid Fixed Area Data
20-23	Machine Check Old PSW (48-55) 20 AMWP 21 Masks and Protect Key 22 Program Mask and Condition Code 23 Instruction Address
27	Floating Point Registers (352-383)
28	General Purpose Registers (384-447)
29	Control Registers (448-511)
30	CPU Extended Logout
31	Storage (Validity of storage being processed by instructions when interrupt occurred.)

●Figure 50.10.4. Machine check code - Model 165

Page of GC20-1730-0
Revised 7/14/70
By TNL GN20-2227

- External Damage - This interrupt occurs if both PSW bit 13 and the external mask bit are on. The ED bit of the stored machine check code (bit 5) will be on to indicate that the storage control unit or processor storage is damaged but that operations associated with the CPU were not affected because the multiple-bit storage failure was associated with an I/O operation. When this type of error occurs, both a soft machine check and an I/O interrupt occur. The error is logged after the machine check interrupt. Error recovery procedures are initiated after the I/O interrupt.

- Time of Day Clock Damage - This interrupt occurs if both PSW bit 13 and the external mask bit are on. The CD bit in the stored machine check code (bit 4) will be on to indicate that an error occurred in the time of day clock that renders the clock invalid. Once this invalid indication has been given, subsequent STORE CLOCK instructions cause the condition code in the current PSW to indicate the fact that the clock is invalid. Error logging is required as a result of clock failure.

Hard machine check interrupts on the System/370 Model 165 are as follows:

- System Damage - This interrupt occurs if PSW bit 13 is on. The SD bit in the stored machine check interrupt code (bit 0) is used to indicate that an error occurred during the execution of the instruction indicated by the machine check old PSW. The error was either a multiple-bit processor storage failure, a storage protection failure, a CPU error that was unretryable, or a CPU error that could not be corrected by the CPU retry hardware.

  If a multiple-bit processor storage failure caused the interrupt, the SE bit in the stored machine check code (bit 16) is on also, and the address of the failing storage area is indicated in the extended logout area. A storage protection failure is indicated by the KE bit (bit 18).

  A logout to the CPU extended logout area also occurs. For a CPU error, an extended logout occurs on the first and seventh retry. Error logging and the execution of recovery procedures are required after this type of interrupt.

Two modes of system operation will be used: full recording mode and quiet, or nonrecording, mode. In full recording mode, all machine check interrupt types cause an interrupt to be taken and logouts to occur. This will be the normal mode of Model 165 operation. In quiet mode, all or certain soft machine check interrupts are disabled. Quiet mode can be used to permit system operation without error recording for all or certain soft errors when a large number of transient (correctable) errors are occurring. It can also be used to allow Model 165 operation under the control of an operating system without Model 165 machine check handling routines included.

A hard stop status and a hard stop bit have been defined for the Model 165. The hard stop bit is located in control register 14 with the other mask bits discussed. If a hard stop condition occurs, the Model 165 system ceases all operations immediately without the occurrence of a logout to the fixed area. Hard stop is initiated by hardware rather than by programming.

When enabled, a hard stop occurs if a System Damage hard machine check type of error is detected when all machine check interrupts are disabled or during the processing of a previous machine check error.

Implementation of a hard stop prevents system operations from continuing when the nature of the machine malfunction prevents the system from presenting meaningful status data.

The state of the Model 165 after IPL or a system reset is:

1. CPU retry and an extended logout are disabled when a CPU error occurs. A hard machine check interrupt results on any CPU error, including those caused by a buffer malfunction.

2. Recovery reports are disabled. Therefore, single-bit processor storage error corrections do not cause a soft machine check interrupt.

3. External damage reports are enabled. A double- or multiple-bit processor storage failure associated with an I/O operation and damage to the time of day clock cause a machine check interrupt.

4. PSW bit 13 is normally enabled by the IPL PSW (it is disabled by system reset) so that System Damage machine checks (unretryable CPU failures, unsuccessfully retried CPU errors, storage protection failures, and multiple-bit processor storage errors associated with the CPU) cause a hard machine check interrupt.

5. Hard stop is enabled.

6. CPU extended logout is enabled and control register 15 points to location 512 as the beginning of the CPU extended logout area.

7. Channel logouts are disabled.

A machine check situation in a Model 65 or 75 results from hardware detection of a machine malfunction or of a parity error. Bad parity can occur in main storage, in local storage, in a register, in an adder, etc. Error correction hardware is not included in these two models. If the machine check mask in the current PSW (bit 13) is enabled, a machine check on Models 65 and 75 causes an interrupt and a diagnostic scan-out occurs, starting at location 128.

If recovery management support (RMS) is included in an OS control program for the Model 65, the MCH routine gains CPU control to record the error and to attempt corrective procedures after the machine check interrupt. Model 65 MCH performs programmed instruction retry for certain instructions only. Alternately, if an SER routine is present, the error is only logged, since a retry of the failing operation is not provided by this routine.

RECOVERY MANAGEMENT SUPPORT (RMS) - OS MFT AND MVT

MCH and CCH Routines

RMS for the Model 165 consists of extensions to the facilities offered by RMS routines currently provided for Models 65 and up. The two RMS routines, machine check handler (MCH) to handle machine check interrupts and channel check handler (CCH) to handle certain channel errors, will be included automatically in MFT and MVT control programs generated for the Model 165.

The two primary objectives of RMS are (1) to reduce the number of system terminations that result from machine malfunctions and (2) to minimize the impact of such incidents. These objectives are

accomplished by programmed recovery to allow system operations to continue whenever possible and by the recording of system status for both transient (corrected) and permanent (uncorrected) hardware errors.

MACHINE CHECK HANDLER: During IPL of a control program containing Model 165 RMS routines, machine check mask bits are enabled, and control register values are set to permit all machine check interrupts and logouts to occur.

MCH receives control after the occurrence of both soft and hard machine check interrupts. When a soft machine check occurs (successful CPU retry, single-bit processor storage error corrected, time of day clock damage, or multiple-bit processor storage error during an I/O operation), MCH formats a recovery report record to be written in the system error recording data set SYS1.LOGREC. This record contains pertinent information about the error, including pertinent data from the logout areas, an indication of the recovery that occurred, identification of the job, job step, and program involved in the error, the date, and the time of day. The operator is informed of successful CPU retries, single-bit processor storage corrections, and an error in the time of day clock. System operations continue except after a TOD clock error.

MCH performs an additional function when a CPU retry was necessary because of a buffer malfunction. When an error occurs in the buffer, as indicated in the extended logout area, MCH updates a programmed buffer error counter. After a certain number of buffer errors occur, the entire high-speed buffer is disabled and MCH notifies the operator of this fact. The operator can allow the system to continue running in degraded mode, if necessary. All CPU fetches are then made directly to processor storage, bypassing the buffer. Alternately, the operator can terminate system operations and request CE diagnosis and repair of the buffer.

Prior to relinquishing CPU control, MCH determines whether or not an automatic mode switch from recording mode to quiet mode should take place if a CPU retry or an ECC correction recovery has just occurred. The determination of whether to switch to nonrecording (quiet) mode is made on the basis of the number of soft machine checks of a specific type that occur during system operation. Error count thresholds are maintained separately for successful CPU retry and successful processor storage single-bit error corrections. The IBM-supplied threshold values can be altered when the control program is generated.

MCH switches the system to quiet mode for either ECC corrections only (the DIAGNOSE instruction is used to change the ECC mode bit from full recording to quiet mode) or for both CPU retry and ECC corrections (the System Recovery mask bit is disabled). Mode switching occurs if the number of soft machine checks that occur during system operation exceeds the specified error count threshold for that type (or if SYS1.LOGREC is full). The operator is informed of the mode switch and can switch back to recording mode at any time thereafter.

Mode switching is implemented to attempt to prevent SYS1.LOGREC from being filled with recovery reports when a recurring correctable error condition exists that would cause many reports to be generated.

When a System Damage hard machine check occurs (uncorrectable or unretryable CPU error, multiple-bit processor storage error, or a storage protect key failure), MCH determines whether the error is one that is correctable by programming. A multiple-bit processor storage error or a storage protect key failure associated with CPU processing causes control to be given to the repair portion of the program damage assessment and repair (PDAR) routine of MCH. PDAR can repair damaged control program storage areas by loading a new copy of the affected module if the module is marked reentrant and refreshable (it has been

written in a read-only manner that will allow reloading of the module at any time without altering execution results). Only refreshable modules residing in the control program nucleus, the link pack area, or an SVC transient area that are damaged will be repaired, if possible. Processing program modules will not be refreshed.

If PDAR cannot correct the error or if the error is an uncorrectable type, PDAR attempts to identify the task associated with the error so that the task can be terminated abnormally. A damage report record, which contains both the fixed and CPU extended logout area data, the recovery action taken, the program identification, the date, and the time of day, is prepared and logged.

System operation continues if the error was corrected or if the error task associated with an uncorrectable error can be identified and terminated. System operation halts, and a re-IPL is required if an uncorrectable error damages a portion of the control program or if the error cannot be associated with a specific task. (A special emergency recording routine attempts to record any existing machine check records in SYS1.LOGREC.) The operator is informed of whatever action is taken.

MCH for the Model 165 contains model-dependent routines and will not execute correctly on System/360 models or on another System/370 model. (See Section 60:15 for a discussion of operating system portability.)

CHANNEL CHECK HANDLER: CCH receives control after a channel error causes an I/O interrupt. CCH formats both an error information block (containing channel logout data) for use by an ERP routine and a CCH error record for recording in SYS1.LOGREC. The latter contains status information from the logout area, program identification, date, and time of day.

If CCH determines that operating system integrity has been impaired by the channel error, control is given to the emergency MCH routine for error recording, and system operations are terminated. Otherwise, the error information block and error record are passed to the appropriate device-dependent error recovery procedure (ERP), which logs the error record and retries the failing I/O operation, if possible, using status information from the error information block. If a successful retry occurs, system operation continues. If the error is deemed permanent (uncorrectable), another error record is prepared and recorded by the outboard recorder routine (OBR), and the task involved is abnormally terminated (unless I/O RMS or a user-written permanent error-handling routine is present). The operator is informed of the abnormal termination and system operation continues.

The CCH routine is structured in a manner that makes it model independent. A channel/model-independent module resides in the operating system nucleus. The required channel-dependent modules included in the operating system at system generation time are loaded during the IPL procedure by NIP (nucleus initialization program), which determines the types of channels that are present in the system, using the new STORE CHANNEL ID instruction and channel configuration data specified by the user at system generation time. CCH routines are therefore compatible for System/370 Models 155 and 165 and for System/360 Models 65, 75, 85, 91, 195, and MP/65 systems.

ERROR RECOVERY PROCEDURES (ERP's)

These device-dependent error routines are a standard part of the control program generated for any OS environment. They will accept and use the Model 165 channel logout data formatted by the CCH routine.

When a channel or an I/O device error occurs on a Model 165, the appropriate ERP is scheduled to perform recovery procedures. If the error is corrected, operations continue normally. If the error cannot be corrected (it is permanent), error recording occurs (see OBR/SDR discussion below). If I/O RMS or a user-written permanent error-handling routine is not present, the affected OS task is abnormally terminated. The operator is notified of permanent I/O errors. ERP routines written for the 3330 and 2305 facilities and the 3211 Printer will include support of the larger number of sense bytes provided by the control units of these devices.

STATISTICAL DATA RECORDER (SDR) AND OUTBOARD RECORDER (OBR)

OBR and SDR routines are included in all OS control programs. These routines are requested by the ERP routines during their processing. The SDR routine is requested when one of the error statistics counters becomes full. Counters are maintained in the resident control program storage area for each I/O device in the system configuration. SDR records these statistics in the appropriate SDR summary record for that device contained in the error log data set SYS1.LOGREC. This insures recording of temporary I/O device error data. The OBR routine records both temporary and permanent channel errors and writes an outboard record containing pertinent status data whenever a permanent error occurs for a device. SDR is also executed when a permanent error occurs to write accumulated statistics for that device.

ENVIRONMENT RECORDING, EDIT, AND PRINT PROGRAM (EREP)

OS EREP is a standard system utility that can be initiated as a job step via standard job control statements at any time. It contains model-dependent routines and will be extended to handle all the status records written by OS RMS routines for System/370. It performs the following:

1. Edits and prints all error records contained in SYS1.LOGREC. These records have been constructed and written by MCH, CCH, OBR, and SDR routines.

2. Accumulates a history of specified record types from SYS1.LOGREC by creating or updating an accumulation data set

3. Edits and prints a summary of selected records from SYS1.LOGREC or an accumulation data set

I/O RMS (APR AND DDR)

I/O RMS routines are optional, model-independent routines supported for MFT and MVT environments. These reconfiguration procedures attempt to minimize the number of abnormal job terminations and unscheduled system halts that occur because of errors on channels or I/O devices.

The alternate path retry (APR) routine provides for the retry of a failing I/O operation on another channel path to the device involved, if one is available, when an uncorrectable channel error occurs. Thus APR, if present, is entered from a device-dependent ERP when a permanent error is deemed to exist after retry procedures have been attempted. If the I/O error is corrected by a retry using the alternate channel path, operations continue. If a permanent error still exists, the task is abnormally terminated unless the DDR routine is present. A malfunctioning channel path can be varied offline by the operator if necessary.

The dynamic device reconfiguration (DDR) routine permits the operator to move a demountable volume from one device to another of the same type when a permanent hardware error occurs and provides repositioning of the volume so that the failing I/O operation can be retried. A volume can also be demounted, so that device cleaning procedures can be performed, and then remounted on the same device. The DDR option also supports demountable system residence devices and unit record equipment. DDR is entered from a device-dependent ERP after a permanent error, either channel or device, occurs on a demountable device. Task termination occurs if the error cannot be corrected and a user-written permanent error-handling routine is not present.

ADVANCED CHECKPOINT/RESTART AND WARM START FACILITIES

If the RMS and I/O RMS routines fail in their attempt to correct a hardware error and the error is one that causes a program or system termination, the automatically provided advanced checkpoint/restart and warm start facilities of OS can be employed to minimize the impact of these terminations on system operation. The automatic restart facility can be used to cause terminated programs to be rescheduled immediately without resubmission of their job control, so that a minimum of operator intervention is required. The operator must authorize all automatic job step restarts. If a permanent I/O device or channel failure caused the program termination, the device or channel can be varied offline. This will insure allocation of a different device when the program step is reinitiated.

The warm start facilities of the control program provide automatic saving of SYSIN and SYSOUT data sets and input and output work queues so that processed work is not lost when a system termination occurs. The operator is informed of the status of jobs in execution when the system terminated, and these jobs should be restarted automatically from the beginning or from a checkpoint if the type of processing involved permits such a restart. System design should include planned restart procedures for unscheduled terminations of individual programs and the system.

50:15 REPAIR FEATURES

The programmed repair features supplied are designed to minimize the impact of malfunction diagnosis and repair on system availability. Fault location and repair time should be reduced by:

1. Improved error recording. Both intermittent and solid hardware and I/O errors are logged at the time of failure if Model 165 RMS routines are present. More status information will be available than was recorded previously for Models 65 and 75.

2. Online error diagnosis. Error diagnosis and repair can be performed concurrently with system operation in a multiprogramming environment to avoid total system or direct access facility unavailability.

3. CE diagnostic routines. More exacting diagnostic routines will be available.

The following maintenance and diagnostic routines will be provided:

- Online Test Executive program (OLTEP) and Online Tests (OLT's) for operation under OS control to test malfunctioning I/O devices concurrently with system operation. OLT's are provided for most Model 165 I/O devices.

- Processor Logout Analysis Program for operation under OS to analyze SYS1.LOGREC MCH records and attempt to identify replaceable malfunctioning hardware components.

- System Test, Channel Test, CPU Test, and Storage Test programs for customer engineer use in identifying malfunctioning hardware units in the system (stand-alone routines).

- Microdiagnostics for customer engineer use in diagnosing malfunctioning field-replaceable units (stand-alone routines).

OLTEP AND OLT's

OLTEP is designed to operate as a processing program under operating system control and is a system generation option. The OS OLTEP provided for operation on System/370 contains functions in addition to those provided by the currently available OS OLTEP. OLT's must be ordered from the customer engineer, link-edited, and placed in a library (PDS) for use under OS OLTEP control. A stand-alone version of OLTEP, called OLTSEP, will be available as well.

OLTEP directs the selection, loading, and execution of device-dependent OLT's for the purpose I/O device testing and error diagnosis. OLTEP provides the required interface between the operating system and the device-dependent OLT's.

As with any other job step, OLTEP is invoked with job control and executes with a user-assigned priority. A minimum program area of 16K is required for OLTEP operation in OS environments. The input stream or system console device can be used to supply the parameters required for test operations - devices to be tested, options desired, etc. (The ability to use other than a console is a new feature.) OLTEP insures the protection and security of user data sets and processor storage in use while OLT's are operating. It also insures that the devices to be tested are online or offline (as far as the operating system is concerned), as required by the particular device type.

OLTEP has the new capability of being able to access the error log data set (SYS1.LOGREC) to obtain history records describing previous I/O errors on the device being tested. In addition, multiple devices can be tested during one OLTEP execution. If a console is used to define the test run, prompting can be requested as an aid to the operator supplying the definition.

OLTEP and OLT's operate concurrently with other executing jobs in a multiprogramming environment and provide online channel, control unit, and I/O device testing and repair, eliminating the necessity for complete system unavailability.

PROCESSOR LOGOUT ANALYSIS PROGRAM

This routine runs under OS control and performs the following:

1. Formats and prints the machine check records written in SYS1.LOGREC

2. Creates and maintains a logout history on a secondary data set

3. Analyzes the MCH logout data records that describe hardware errors and prints the location of suspected malfunctioning replaceable units. The history logout data set can be used

in this analysis in an attempt to gain better resolution of the malfunctioning replaceable units.

Since the Processor Logout Analysis program operates under OS control, it can be run concurrently with other operating jobs in a multiprogramming environment and should be run by the operator when the error log becomes full. It will be used by the customer engineer during scheduled preventive maintenance periods to locate units responsible for intermittent hardware failures.

SYSTEM TEST, CHANNEL TEST, CPU TEST, AND STORAGE TEST PROGRAMS

These programs are stand-alone, model-dependent, diagnostic routines for customer engineer use. The System Test performs a functional test of the hardware system components, including the CPU, the buffer, storage, all channels, and most locally attached I/O devices. The 2701 Data Adapter Unit and the 2702 and 2703 Transmission Control Units are supported in wrap or diagnostic mode only (no outboard operations to remote terminals). The Storage Test performs a functional test of processor storage, buffer storage, and the buffer directory. The Channel Test performs a detailed diagnosis of 2870, 2860, and 2880 channels, while the CPU Test analyzes the CPU in more detail than the System Test.

The System Test can be used for the following purposes:

1. To locate a malfunctioning unit in the system. (The specific error component within the unit is not identified.)

2. To determine a working hardware configuration prior to a system IPL

3. To verify correct system operation after a malfunctioning unit has been repaired

These diagnostic tests are designed for simplicity of operation. They are loaded from cards, tape, or disk and use the system console for communication purposes. Users can run the System Test at appropriate times, for example, prior to IPLing the system for an extended period of online operation.

MICRODIAGNOSTICS

Microdiagnostics are a set of stand-alone, high-resolution, fault-locating tests. A set of resident microdiagnostics is contained on a magnetic disk cartridge. They are loaded into WCS from the device in the system console. The resident microdiagnostics are designed to test those parts of the CPU required to load WCS with nonresident microdiagnostics that are contained on an I/O device (card reader, tape drive, direct access device). Nonresident microdiagnostics are to be used by customer engineers to locate malfunctioning field-replaceable units when a solid hardware error exists.

50:20 RAS SUMMARY

The degree to which an installation benefits from available RAS features depends in part upon their proper implementation. It is desirable for Model 165 users to design a system that includes RAS features and to become involved in the implementation and use of maintenance procedures. Specifically, the user can:

- Include OLTEP and OLT's in his operating system

- Train operators to run the System Test and the Processor Logout Analysis programs. The conditions under which these programs are to be run should be determined jointly by the user and the customer engineer.

- Plan system and program recovery procedures (use of warm start and checkpoint/restart facilities)

- Have operating personnel perform normal hardware maintenance procedures, such as the periodic cleaning of tape drive heads. Proper system hygiene should be maintained, in general.

- Implement an effective program of operator training so that the number of system malfunctions that occur because of operator error is reduced.

SECTION 60: PROGRAMMING SYSTEMS PREINSTALLATION PLANNING

Because extensive hardware and programming systems compatibility exists between the System/370 Model 165 and System/360 models, most Model 65 and 75 OS users can upgrade to a Model 165 with a minimum of effort. This is also true for users of ASP (Attached Support Processor) who wish to upgrade a Model 65 main processor to a Model 165. Essentially, no more effort may be involved in the installation process for OS MFT and MVT users than is required currently to change from one operating system release to another, or to regenerate an operating system to include new hardware, new I/O devices, and more control program options.

It also may be desirable to expand the design of the system - add new applications, use additional control program facilities, etc. - as a result of a sizable upgrade in the hardware configuration. In most cases, the fact that an MFT or MVT user is upgrading to a Model 165 should not add to the effort that would be required if these system changes were to be implemented for a Model 65 or 75 upgrade to another System/360.

60:05 OS MFT AND MVT TRANSITION

A system generation must be performed using an OS release that includes Model 165 support in order to obtain an MFT or MVT operating system designed to support new Model 165 features. The existing system generation job stream can be used with the following modifications, as appropriate:

- Direct access space allocation for operating system data sets will have to be adjusted as indicated in IBM System/360 Operating System, System Generation (GC28-6554). If a 3330 or a 2305 facility is to be used as a system residence device, UNIT parameters in job control statements must be changed where necessary.

- Stage I input must be modified to reflect the Model 165 configuration, including the presence of any new I/O devices or features, such as the 3211 Printer, the 3330 and 2305 facilities, an integrated emulator program, etc. Other control program options, such as I/O RMS, OLTEP, and performance improvement features, can be included. MCS and DIDOCS must be requested to support the CRT and keyboard as an operator's console.

- FCB and UCB images should be added to SYS1.SVCLIB if a 3211 Printer is included in the configuration. User-written output writer procedures should be modified to include these specifications.

The OS MFT or MVT operating system generated for the Model 165 includes the following:

- A nucleus designed to operate in the fixed storage area of the Model 165. RMS routines (MCH and CCH) that support the expanded machine check interrupt are included. OBR and SDR are present also.

- Control program support of block multiplexing, and rotational position sensing as discussed in Section 30, if requested

- Support of the user-specified new I/O devices and Model 165 operator console

- Support of the interval timer and time of day clock
- Support of the new instructions by Assembler F
- The required interface to the integrated emulator specified at system generation, if any

If integrated emulator programs are to be used, the steps outlined in Section 40 must be taken in order to convert from current emulation procedures to Model 165 emulation.

EXISTING PROCESSING PROGRAMS AND JOB CONTROL

IBM-supplied OS program products and Type I processing programs (language translators, utilities, etc.) will run on the System/370 Model 165 without alteration. Subject to the exceptions stated in Section 10:05, user-written OS processing programs that operate on Models 65 and 75 will also execute correctly on the Model 165. Modification and reassembly of existing user-written OS processing programs is not required unless new processing is to be added or existing processing is to be altered.

Modification of the job control for these processing programs is required if I/O device type is changed, say from 1403 to 3211, if direct access space allocation changes, if a DCB parameter is to be altered, etc. I/O device type changes do not necessitate processing program alterations unless device-dependent macros have been used, data organization is changed, or a DCB parameter specified in the program is to be altered.

CONVERSION TO 3330 AND 2305 FACILITIES

Conversion from current direct access devices to the 3330 facility involves the same procedures as are required now to change from one disk device to another, say from 2311's to a 2314. Existing disk data sets can be placed on 3336 Disk Packs by using an IBM-supplied utility in most cases. Assuming that data organization is not changed, consideration should be given to altering the block size used and the amount of space allocated to the data set. The location and size of each type area in an ISAM data set should be altered, taking into account 3336 Disk Pack characteristics. These changes can be made in job control statements.

Sequentially organized data sets (processed by QSAM or BSAM) and partitioned data sets can be copied from the source direct access device directly to the 3330 facility by use of the OS IEHMOVE utility. Or they can be copied to tape and then to the 3330 facility by use of the same utility (if the source direct access device type is not present in the Model 165 configuration).

Indexed sequential data sets can be copied directly from the source direct access device to the 3330 facility by use of the IEBISAM utility. Alternately, they can be unloaded to tape and then reloaded by means of the same utility. Changes to space allocation, etc., can be made via job control statements.

Direct organization (BDAM) data sets can be copied on a track-to-track basis from the source direct access device to the 3330 facility by using IEHMOVE or copied to tape and then to the 3330 facility. If more records are to be placed on a 3330 track than are on a source device track, the existing reorganization program can be used to transfer the data to the 3330 facility, and the program may have to be changed. Reprogramming of the randomizing routine used in the

reorganization, and all processing programs that access the BDAM data set, is necessary if a relative track or actual address reference is used and fewer (or more) 3330 tracks are allocated to the data set than before.

Subject to the restrictions stated in Section 10:05 and those indicated for BDAM data sets, existing executable processing programs can be used without change to process data sets that have been transferred to 3336 packs. Nothing need be done to job control for these programs if the cataloged procedures supplied with the language translators are used, as long as the 3330 facility is specified as a SYSDA device at system generation. Otherwise, job control statements must be changed to request 3330 devices and, optionally, any data set characteristic changes, such as block size. RPS support, as described previously, is provided automatically.

User-written programs that use the EXCP level to access disk data sets that have been transferred to 3336 packs may have to be modified to reflect the characteristics of the data set on the 3336: a different number of records per track, a different number of tracks per cylinder, etc. All 2311 and 2314 CCW lists will operate on 3330 facilities except those that are device or channel time dependent and those that use the file scan commands, which are not available on the 3330. User-written 2311 or 2314 error routines will not execute correctly and must be modified. RPS commands have to be added by the user if this support is desired for programs that use EXCP. (Note that the XDAP macro will include support of RPS commands.)

Data sets currently located on 2301 Drum Storage can be placed on 2305 modules by means of a data set utility program. Unit specification in the job control statements of existing programs that will access the 2305 module instead of the 2301 drum must be changed. Also, to reflect the smaller capacity of a 2305 track, it may be necessary to alter the block size used. The latter can also be done via job control statement alterations (without program reassembly) unless block size was specified in the program itself or record length is larger than 2305 track size.

CONVERSION TO THE 3803/3420 MAGNETIC TAPE SUBSYSTEM

As stated in Section 20:25, existing tape processing programs and their job control statements and tape volumes need not be modified in order to be used with 3803/3420 subsystems with equivalent features whenever the same recording modes are used.

Whenever possible, seven- and nine-track NRZI mode tape volumes should be converted to 1600-BPI PE format to obtain the benefits of the higher density and the PE recording technique. In cases in which tape volumes must retain seven-track NRZI format, for interchange with other systems for example, use of the 3803/3420 subsystem offers improved tape reliability and subsystem serviceability as already discussed.

Conversion of seven- and nine-track NRZI tape volumes can be done gradually during production processing. That is, the old master input tape volume is read in on a 3420 tape unit with the appropriate compatibility (Dual Density or Seven-Track) feature while the new master output tape is written on a 3420 tape unit in 1600-BPI PE format. Existing programs that process these converted tapes need not be modified unless an altered characteristic (recording mode or block size, for example) is specified in the program DCB. Existing job control for these programs must be altered to request a tape unit with the new recording characteristics and, if desired, to change existing DCB parameters such as block size, number of buffers, etc.

Tapes that cannot be converted on an as-used basis, such as program tapes or active reference tapes that are not rewritten when processed, can be converted by using a copy utility.

60:10 PLANNING OPTIMAL SYSTEM PERFORMANCE, USING BLOCK MULTIPLEXER CHANNELS AND RPS DEVICES

Block multiplexing, rotational position sensing, and multiple requesting provide the user with another tool that can improve total system throughput. However, the effectiveness of this tool for a given installation depends largely on proper planning for its use. This section indicates how to use block multiplexer channels and RPS devices effectively.

The guidelines outlined indicate how best to configure a system with rotational position sensing devices, how to plan job scheduling, and what to consider when determining disk data set characteristics. Explanations follow the statement of each guideline.

All guidelines presented are not necessarily practical for all users. Each item should be evaluated in terms of the processing requirements and hardware configuration of an installation.

SYSTEM CONFIGURATION AND GENERATION

System configuration and generation guidelines are as follows:

1. Multiple 3330 facilities should be placed on a single block multiplexer channel.

   Performance improvement occurs (1) as a result of overlapping the rotational positioning time of disk devices and (2) because more I/O requests can be initiated in a given period of time, since the channel is free more often. When many disk devices are active concurrently on a block multiplexer channel, there is more potential for such overlap.

2. Direct access devices with RPS should be placed on separate channels from I/O devices without RPS. Alternate possibilities are as follows:

   a. If it is necessary to place non-RPS devices on the same block multiplexer channel with RPS devices, give first choice to non-RPS devices with a buffered control unit, such as the 2540 Card Read Punch and the 1403 Printer. These devices disconnect from a block multiplexer channel during the relatively long mechanical portion of their cycle, thereby freeing the channel for other operations.

   b. Tape units should not be placed on a block multiplexer channel with RPS devices unless absolutely necessary, because channel disconnection does not occur during any of their channel operations. If this is not possible, try to plan job scheduling to avoid having jobs using tape units and jobs using RPS support active on a block multiplexer channel at the same time. If this is not feasible, try to assign very low-activity data sets to these tape units.

      A device without channel disconnect capability can monopolize the block multiplexer channel for relatively long periods of time, thereby preventing (1) the initiation of other I/O operations on the channel and (2) the reconnection and completion of disk RPS channel programs already in operation

on the channel. For example, a direct access device without RPS retains use of the channel during its search operations as well as during its reads and writes. If the device is a 2314 and block size is half a track, the channel is busy for 25 ms on the average (12.5 ms average rotational delay plus 12.5 ms read/write) for each I/O operation started for the non-RPS 2314 facility. Even if the block size used is relatively small, the channel can still be monopolized by the non-RPS device if there is high activity on the device.

3. The 2305 facility normally should not be placed on a block multiplexer channel with other devices.

   Exclusive use of a channel insures optimum performance of the 2305 facility as a system residence device.

The following should be noted in regard to specification of priority and ordered-seek I/O request queuing options for RPS devices at OS system generation. The priority queuing option insures priority I/O request initiation for the device, but because of first-come, first-served handling of I/O operations on the block multiplexer channel, this option does not insure that priority device channel programs will complete sooner than other RPS channel programs that were started later on the channel. However, the objective of specifying the ordered-seek queuing option (minimization of arm movement on a disk drive) can still be achieved when RPS is used.

THIS PAGE LEFT INTENTIONALLY BLANK

JOB SCHEDULING

The following are guidelines for job scheduling.

1. If total system throughput improvement, rather than maximum individual job performance increase, is the objective, schedule jobs that use RPS together, such that the maximum number of RPS devices are active concurrently on each block multiplexer channel.

   The more RPS devices active concurrently on a block multiplexer channel, the greater the overlap potential that exists. (See item 1 under "System Configuration and Generation".)

2. When jobs that use disk data management functions with RPS support are executed concurrently with jobs that do not use RPS support, the devices assigned to the former should be on different channels from devices assigned to the latter.

   If jobs using RPS and jobs not using RPS must access devices on the same block multiplexer channel concurrently, the jobs without RPS support should have high seek activity such that search and read/write time is small compared to seek time.

   Assume sequentially organized data sets and TCAM message queues are allocated to the same 3330 facility on a block multiplexer channel. A QSAM job step and a TCAM job step are executing and access the 3330 facility concurrently. Since RPS commands are not used for TCAM message queue processing, each 3330 disk drive containing TCAM queues acts like a non-RPS device and can monopolize channel time. Thus, the job steps that use QSAM can be delayed. (See item 2 under "System Configuration and Generation".)

3. Allocate a data set that will be accessed using QSAM or BSAM chained scheduling to a device on a channel without active RPS jobs.

   The chained scheduling technique is designed to keep a device active as long as record processing keeps up with record reading or writing. Thus, either a selector or a block multiplexer channel can be kept busy for long durations, preventing the execution of any other I/O operation on the channel. Note that while QSAM and BSAM support concurrent use of RPS and chained scheduling for access to a disk data set, the performance attained by using chained scheduling alone will not be improved significantly by using RPS as well.

4. When data sets are being processed by an RPS access method in a multiprogramming or multitasking environment and disk device assignment is handled by the user rather than by the control program, allocate as many separate RPS devices as is practical.

   This approach allows the possibility of having more concurrent operations on these data sets and therefore more seek and rotational positioning overlap.

5. If a response-oriented RPS job operates on a block multiplexer channel concurrently with other RPS jobs, job scheduling should insure that the number of jobs executing simultaneously is such that the performance desired for the response-oriented job can be attained.

   The performance of a block multiplexer channel is affected by the percentage of time the channel is busy searching and reading.

The read or write of a particular record may be delayed one rotation because the channel is busy servicing another channel program. The probability of a particular record being delayed is a function of the percentage of channel busy time. As block multiplexer channel utilization increases, the probability that individual channel programs will be delayed increases. It is theoretically possible for the read or write of a particular record to be delayed indefinitely because the block multiplexer channel is busy searching for and reading other records. That is, utilization of more and more block multiplexer channel time will normally result in better overall performance but will increase the likelihood of delayed response from any one data set.

DATA MANAGEMENT PARAMETERS

Following are guidelines for establishing data management parameters:

1. When organizing direct data sets to be processed using BDAM, use fixed-length standard or VBS records and a record reference that includes ID (relative block, relative track and ID, or actual address).

   RPS is supported only for fixed-length standard and VBS formats without key reference because record position must be known in order to calculate the sector number required for positioning. However, if a key reference or a variable record format is used, RPS support is provided for write verification and update (after retrieval) operations.

2. Use a large block size for sequentially processed data sets whenever possible, subject to the availability of processor storage for buffers.

   The use of RPS can provide performance gains for both short and long disk record blocks. However, use of large blocks rather than short reduces the total time required to read or write a given data set, because less disk space is required and fewer I/O operations are necessary. Note also that total throughput for a given block multiplexer channel is improved by using blocks of equal (or nearly equal) size for all data sets being processed on the channel.

3. Use fixed standard records for QSAM and BSAM data sets where possible.

   The channel programs used for fixed standard records free the channel more often than when other record formats are used. A search for the previously read record in order to locate the next sequential record is not used when fixed standard records are read sequentially. The sector number of the next sequential record is obtained by including a READ SECTOR command at the end of the channel program used to read each record. Therefore, the SEARCH command specifies the ID of the desired record, and the channel is free during the time it would otherwise have been busy searching for the previously read record.

   Channel time is reduced when standard records are written, or because the operation required to calculate the remaining number of bytes on a track after each write is eliminated. (Note that the disk control unit is still busy erasing to end of track after formatting write operations even though the channel is freed after the data record has been written.)

4. Use multiple buffers with QSAM and BSAM.

   The availability of multiple buffers per data set lowers the probability that a task will have to wait for a particular record. QSAM is designed to initiate an I/O request whenever a buffer becomes available, thus keeping the channel queue as full as possible. When BSAM is used, the programmer must handle the initiation of I/O requests.

The following summarizes the advantages of rotational position sensing, multiple requesting, and block multiplexing:

- System throughput increases can be achieved when multiple sequential data sets are processed concurrently on a single block multiplexer channel (using QSAM, QISAM, or BSAM) because a higher effective channel data rate results.

- The number of block multiplexer channels required in a given system configuration can be fewer than the number of selector channels that would be required to handle the same amount of data, because more effective channel utilization is achieved by block multiplexing disk operations.

- The performance cost to an installation of verifying disk write operations is sharply reduced.

- The greatest throughput improvement results from use of rotational position sensing with high-activity, transaction-based processing, that is, with applications that include one or more large jobs that:

  1. Use direct processing (BDAM) with fixed-length standard records and a record reference that includes ID

  2. Require a multivolume data base of small records

  3. Process many additions and updates and use write verify

60:15 OS PORTABILITY

To avoid multiple system generations an OS user with multiple Model 165 systems may wish to generate a single operating system that can be used on every Model 165 in the installation. This is possible under the same system hardware and I/O device configuration restraints that exist for System/360 models. That is, during the IPL procedure, channels and I/O devices may have to be varied offline, partition sizes may have to be redefined, etc., when the operating system is used with a different configuration than was specified during system generation.

A user with both a System/370 Model 165 and a System/370 Model 155 or a System/360 model in an installation may also wish to generate one operating system that can be used on both models. This approach provides backup when one system is unavailable and can eliminate the necessity of multiple generations.

Portability of an OS operating system between a System/370 Model 165 and a System/370 Model 155 or a System/360 model, say 50 or 65, can be achieved by utilizing a multiple nucleus control program under the following general conditions:

1. The system hardware and I/O device configuration of both systems must be similar. For example, a Model 165 OS control program generated to support block multiplexing mode and RPS direct

access devices cannot be executed on a system without such channels and devices.

2. The same control program, MFT or MVT, must be used for both systems.

3. Consideration should be given to the processor storage sizes of the two models when determining such factors as size of the scheduler, language translators, and linkage editor(s) included in the generated system.

4. Processing programs that are to run on both models must use instructions and features common to both systems. An Assembler Language program that uses the new general purpose instructions for the Model 165 or byte orientation can be executed on a Model 155 but not on Models 50 or 65.

In order to generate an operating system that is portable between the Model 165 and the Model 155, 50, or 65, the following steps are required:

1. A complete system generation must be performed to generate an operating system for the Model 165. The IPL-time system/operator communication option must be requested so that options specified can be altered during IPL.

2. A nucleus generation should then be done for the alternate system. The model number specified (in the SUPRVSOR, SECMODS, CENPROCS macros, etc.) will be that of the alternate system, not the Model 165.

3. Additional link-edits must be performed to add model-dependent routines to the generated multiple nucleus operating system. Specifically, MCH or SER and EREP model-dependent routines for the secondary system must be included, as appropriate.

4. If extended precision floating-point divide is used in processing programs, the following steps should be taken. SYS1.LINKLIB contains two divide simulation routines. One uses extended precision hardware, the other does not. When a full generation is performed for the Model 165, a calling mechanism is set to request the divide routine that uses extended precision instructions at execution time, since the Model 165 contains these instructions.

   Therefore, the divide simulation routine that does not use extended precision should be transferred from SYS1.LINKLIB to another library and given the same member name as the divide routine with extended precision instructions. When the operating system is executed on a Model 165, SYS1.LINKLIB should be used by extended precision programs. When the operating system is executed on a Model 50 or 65, the alternate library should be used.

Whenever a new program that is to be used on both systems is added to a library or if the Model 165 hardware configuration changes, the user must consider whether or not portability is affected.

60:20 USE OF OTHER PROGRAMMING SYSTEMS

Subject to the restrictions stated in Section 10:05, users of OS PCP, DOS, TOS, BOS, and OS MFT and MVT control programs not generated for a Model 165 or non-IBM-supplied control programs can execute their existing control and processing programs on a Model 165 with a hardware

configuration comparable to that of the current System/360 model. (Certain BPS programs have known timing dependencies that prevent their successful execution on a Model 165.)

Note that a control program not generated specifically for the Model 165 does not include support of the Model 165 CRT and keyboard as the primary console device. Therefore, such a control program cannot be used on the Model 165 unless it supports a device as a console that is also present on the Model 165. The console device must have the same I/O device address on the Model 165 as it has on the current system. In addition, Model 165 RMS (machine and channel check routines) are not included in these programming systems, and the Model 165 operates under the conditions listed below.

- Single-bit processor storage errors will be corrected by ECC hardware, but IPL disables this interrupt (recovery bit masked).
- CPU retry will not occur, since it is disabled at IPL. Any CPU error, including those caused by a buffer failure, will result in a hard machine check condition and generation of CPU extended logout data.
- External Damage and Time of Day Clock Damage are enabled and cause a machine check condition.
- Any hard machine check error (an uncorrectable or unretryable CPU error, a double- or multiple-bit processor storage error, or a storage protection failure) will cause a hard machine check condition and generation of CPU extended logout data.

A machine check control switch that determines what action is taken when a machine check condition occurs is present on the system console. When this switch is set to the PROCESS position, machine checks cause an interrupt and a logout if they are not disabled. This setting is to be used when an operating system containing Model 165 RMS is in operation. When the switch is in the STOP ON CHECK position, all machine checks cause a hard stop without an interrupt or a logout.

If the system is not set to hard stop after a machine check when an operating system without Model 165 RMS is used, the system takes whatever action was planned for machine checks, as follows:

- A control program without a recovery routine included (for example, SER0, SER1, MCH for OS, or MCRR for DOS) normally enters the wait state after a machine check interrupt and logout. The logged data can be obtained with a stand-alone storage dump routine or displayed on the graphic console. The operator can re-IPL and attempt to continue operations or the CE can perform diagnostic procedures.
- A control program that contains a recovery routine will enter the routine and attempt execution. As stated in Section 10:05, these routines access model-dependent data and will not operate properly. In addition, the extended logout data stored when the interrupt occurred will have destroyed the code at locations 512 through 1511. Results are unpredictable. Therefore, the system should be set to hard stop if a recovery routine is present in a control program not generated for the Model 165.

For the following reasons it is advantageous for Model 165 users to install an operating system that includes recovery management support designed specifically for the Model 165:

- The number of re-IPL's necessary because of machine malfunctions can be reduced. Most hardware errors will be corrected either by Model 165 hardware recovery procedures or by RMS routines.

The latter insures the continuation of system operations whenever possible if the error cannot be corrected. This is particularly important during online operations. In Model 165 systems without RMS, hardware errors other than single-bit processor storage errors will necessitate a re-IPL.

- The diagnosis of machine malfunctions by customer engineers will be helped by status information recorded by RMS. This data will be of greatest benefit in diagnosing intermittent errors.

- Versions of OS control programs that include Model 165 RMS routines are the only ones that include support of new Model 165 features, channels, direct access devices, integrated emulators, and the graphic operator console.

SECTION 70: COMPARISON TABLE OF HARDWARE AND OS FEATURES FOR SYSTEM/360 MODEL 65 AND SYSTEM/370 MODEL 165

This table has been included for quick reference. It compares the hardware features and OS support of Models 65 and 165. A dash (-) in a programming system column indicates that the hardware feature does not require programmed support.

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
I. CPU				
A. Internal performance times Model 65	1	2 to 5	-	-
B. Instruction set				
1. Decimal and binary arithmetic	Standard	Standard	All languages except FORTRAN support of decimal arithmetic	Same
2. Floating-point arithmetic	Standard	Standard	All languages except RPG	Same
3. Extended precision floating-point	Not available	Standard	Mnemonics provided by Assemblers F and H. Subroutines in FORTRAN H and PL/I Optimizing	Subroutines and instruction support provided by same languages translators.
4. New general purpose instructions (MOVE LONG, COMPARE LONG, etc.)	Not available	Standard	-	Assembler Language mnemonics for user use
5. High-speed multiply	Not available	Optional	-	-
6. STORE CPU ID, STORE CHANNEL ID, HALT DEVICE instructions	Not available	Standard	-	Used by control program
C. Overlap of instruction fetching and preparation with execution	Instruction unit normally prepares one instruction at a time.	Instruction unit can process several instructions while execution unit executes one instruction.	-	-
	Imprecise interrupts occur only for storage protection violations.	Imprecise interrupts occur on an attempt to store at an invalid processor storage address or a storage protected location.		
D. Interval timer	Standard 16.6 ms resolution	Standard 3.33 ms resolution	Supported	Supported except for time of day requests
E. Time of day clock	Not available	Standard	-	Supported for time of day requests
F. CPU retry (hardware)	No	Standard	-	-

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
G. Machine check interrupt	One level of machine check provided for all machine errors. One machine check mask.	Occurs on both soft (corrected) and hard (uncorrected) errors. Four types are defined. Additional masking is provided.	SER0, SER1, and MCH routines. (See item VI for functions provided.)	MCH routine. (See item IV for functions provided.)
H. Size of fixed storage area in lower core	328 bytes including CPU and channel logouts	1512 bytes, reducible to 512 if the extended logout area is moved	Logout data processed by MCH or SER routine	Logout data processed by MCH
I. Control logic	Microprogram in ROS	Microprogram in both ROS and WCS	-	-
J. CPU cycle time	200 nanoseconds 8-byte data path	80 nanoseconds 8-byte data path	-	-
K. Direct Control	Optional	Standard	Not supported	Not supported
L. Compatibility features (all are optional and mutually exclusive)	1. 7070/7074 2. 7080 (for both 705 and 7080 emulation) 3. 709/7040/7044/7090/7094/7094II	1. 7074 (includes 7070 and 7074) 2. 7080 (for both 705 and 7080) 3. 709/7090/7094/7094II (does not include 704, 7040, or 7044)	Three stand-alone emulators are provided.	Three integrated emulator programs provided, one for each 7000 series compatibility feature
M. Multiprocessor systems	1. Multisystem optional feature permits interconnection of two Model 65s. Main storage is shared (512K or more). Direct control is required.	1. A multisystem feature is not available.	1. OS MP/65 supports Model 65 multiprocessor systems (MVT).	-
	2. The support or main processor in an ASP configuration can be a Model 65. Two or three systems are connected via a Channel to Channel Adapter.	2. A Model 165 can be a a support or main processor in an ASP configuration.	2. PCP, MFT, or MVT can be used for a Model 65.	2. MFT or MVT can be used for a Model 165.

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
II. STORAGE				
A. Processor storage sizes	256K 512K 768K 1024K	- 512K - 1024K 1536K 2048K 3072K	All are supported.	All are supported.
B. Processor storage cycle	750 nanoseconds (for 8 bytes). Two-way interleaving of sequential accesses other than by the channels is provided.	2 microseconds. Storage is 4-way, doubleword interleaved for CPU and channel requests. 32 bytes can be obtained every 2 microseconds.	-	-
C. Processor storage validity checking	Parity checking by byte. No hardware error correction is provided.	ECC checking on a doubleword. Single-bit errors are corrected by hardware.	-	-
D. Processor storage reconfiguration	Not provided for uniprocessor systems	Malfunctioning storage boxes can be configured out of the operational system.		
E. Byte boundary alignment permitted for operands of nonprivileged instructions	No	Standard	-	Programmers can use the hardware facility in Assembler programs.
F. Storage and fetch protection	Standard	Standard	Storage protection is supported.	Same as Model 65
G. 2361 Core Storage	Optional Up to 8 million bytes can be attached.	Cannot be attached	Hierarchy support is provided.	-
H. High-speed buffer storage	Not provided	8K is standard. 8K more can be added	-	-
1. Cycle time		80 nanoseconds	-	-
2. CPU-to-buffer fetch time		160 nanoseconds for 8 bytes; a request can be initiated every 80 ns cycle.	-	-

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
III. CHANNELS				
A. Total number per system	Up to 7	1. Up to 7 2. Up to 12 with Extended Channels special feature	Supported	Up to 12 channels supported
B. 2870 Multiplexer Channel	One or two can be attached.	Same as Model 65	Supported	Supported
C. 2860 Selector Channel (1.3 MB)	A maximum of 6 can be attached.	A maximum of 6 can be attached.	Supported	Supported
D. 2880 Block Multiplexer Channel (1.5 MB). Two-Byte Interface feature permits a 3.0 MB data rate.	Cannot be attached	A maximum of 6 can be attached without the Extended Channels feature, a maximum of 11 with this feature.	-	Supported in selector or block multiplexer mode
E. Aggregate channel data rate	In excess of 4 MB with one 2870 and six 2860s	In excess of 9 MB with one 2870 and eleven 2880 channels		
F. Channel retry data provided	Yes in I/O logout	Yes in I/O logout	Channel retry provided by CCH routine	Channel retry provided by CCH routine
G. Channel-to-Channel Adapter on 2860	Optional	Optional	-	-
IV. OPERATOR CONSOLE DEVICES	1. 1052 Printer-Keyboard (optional) 2. Second 1052 Printer-Keyboard is optional. 3. A 2250 Display Unit and a remote 2150 Console are optional. 4. Other devices can be used as primary and secondary consoles a. 1052 Printer-Keyboard b. 2260 Display Station c. 2250 Display Unit	1. Stand-alone system console unit is standard. It includes: a. A CRT-Keyboard combination for operator/system communication b. An indicator viewer c. A microfiche document viewer d. A processor storage configuration plugboard e. A system activity meter	1. Supported as the primary console 2. MCS option supports multiple consoles. 3. Supported as the primary console 4. Additional consoles supported by MCS and DIDOCS options	1. Supported as the primary system console a. Supported by the MCS options and DIDOCS

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
		f. A device for loading WCS and microdiagnostics 2. Optionally, other devices can be used as secondary consoles as listed for the Model 65.		2. Supported by MCS and DIDOCS options
V. I/O DEVICES				
A. 3211 Printer	Yes	Yes	Supported	Supported
B. 3330 facility	No	Yes on a 2880	-	Supported including RPS and multiple requesting
C. 2305 facility Models 1 and 2	No	Yes on a 2880	-	Supported including RPS and multiple requesting
D. 3803/3420 Magnetic Tape Subsystem	Yes	Yes	Supported	Supported
VI. PROGRAMMED RECOVERY ROUTINES FOR HARDWARE ERRORS				
A. SER0 in control program	Optional (Either SER0, SER1, or MCH must be included.)	-	Logs status data and terminates system operation	Not provided
B. SER1 in control program	Optional (Either SER0, SER1, or MCH must be included.)	-	Logs status data and attempts to terminate affected task and continue system operation	Not provided
C. Recovery management support (RMS) - MCH and CCH in program control	Optional	Required	1. Programmed instruction retry is attempted for some instructions. 2. Provides repair of refreshable control program and processing program modules damaged by a hardware error.	1. Status data logged after both successful and unsuccessful hardware retry of CPU errors 2. Single-bit errors are corrected by ECC hardware. Status data is logged by MCH. Repair of refreshable control program modules damaged by double- and multiple-bit processor storage errors and CPU errors is provided.

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
			3. Programmed storage protect key repair is provided.	3. Programmed storage protect key repair is provided.
			4. Selective task termination if programmed recovery is not possible or fails.	4. Same as Model 65
			5. Error logging of machine check errors in SYS1.LOGREC	5. Both hardware-corrected and uncorrected machine errors are logged in SYS1.LOGREC.
			6. Programmed channel retry after a channel error is possible. Channel errors are logged.	6. Same as Model 65

INDEX

THIS PAGE LEFT INTENTIONALLY BLANK

PREFACE

It is assumed that the reader of this document is familiar with System/360. The reader should have a general knowledge of System/360 architecture, channels, I/O devices, and programming systems support. This guide highlights only those Model 165 hardware, I/O, and programming systems features that are different from those of System/360 models and discusses their significance. Additional, more detailed information regarding System/370 Model 165 hardware and programming systems support can be found in the following SRL publications:

IBM System/370 Model 165 Functional Characteristics (GA22-6935)

IBM System/370 Installation Information, Physical Planning (GA22-6971)

IBM System/370 Principles of Operation (GA22-7000)

IBM System/370 Component Summary:

3830 Storage Control, 3330 Disk Storage (GA26-1592)

IBM System/360 Component Description:

2835 Storage Control and 2305 Fixed Head Storage Module (GA26-1589)

3211 Printer, 3811 Control and 3216 Interchangeable Train Cartridge, Component Description (GA24-3543)

Form-Design Considerations-System Printers (GA24-3488)

Emulating the 7070/7074 on the IBM System/370 Model 165 Using OS/360 (GC27-6948)

Emulating the 709, 7090, 7094, 7094II on the IBM System/370 Model 165 Using OS/360 (GC27-6951)

Emulating the 7080 on the IBM System/370 Model 165 Using OS/360 (GC27-6952)

IBM System/360 Operating System:

Planning for the IBM 3211 Printer, Data Management Macro Instructions and Services (GC21-5008)

CONTENTS

FIGURES

TABLES

512K to 3072K of four-way, doubleword interleaved, two-microsecond processor storage is available - three times the maximum available on the Model 65.

Byte boundary alignment is permitted for the operands of non-privileged instructions to eliminate the necessity of adding padding bytes within records or to blocked records for the purpose of aligning fixed or floating-point data.

Error checking and correction (ECC) hardware that automatically corrects all single-bit processor storage errors and detects all double-bit and most multiple-bit errors is standard.

- I/O devices for the Model 165 are as follows:

  Most currently announced I/O devices for System/360 Models 65 and above can be attached.

  The new 3330 facility is available for attachment to 2880 channels. It offers significantly faster seeks and more than twice the data rate of the 2314 facility, more than three times the capacity of the 2314, and automatic error correction features. The new rotational position sensing and multiple requesting capabilities announced for the 2305 facility are standard.

  The 3330 has an 806 KB data transfer rate, average seek time of 30 ms, and full rotation time of 16.7 ms. Up to 800 million bytes can be contained on an eight drive facility.

  The 2305 facility Models 1 and 2 can be connected to 2880 channels to provide significantly faster data transfer operations.

  The Model 1 has a 3 megabyte data rate, a maximum module capacity of 5.4 million bytes, and average access time of 2.5 ms.

  The Model 2 has a 1.5 megabyte data rate, a maximum module capacity of 11.2 million bytes, and average access time of 5 ms.

  The new high-speed 3211 Printer with a tapeless carriage and an alphameric print speed of 2000 lines per minute is available. The tapeless carriage decreases operator intervention by eliminating carriage tape loading and unloading.

- Extensive hardware and programming systems error recovery and repair features are provided to improve system reliability, availability, and serviceability.

- Compact physical design reduces Model 165 CPU and processor storage space requirements. The Model 165 CPU has three times the number of circuits as a Model 65, in excess of 75,000 versus 25,000, yet a 512K or 1024K Model 165 requires the same amount of space as a 512K Model 65.

As the highlights indicate, Model 65 and 75 users have a broader range of Model 165 configurations to choose from than before when tailoring a growth system with improved throughput and expanded capabilities.

Specifically, the Model 165 offers the following advantages over Models 65 and 75:

Larger Processor (Main) Storage Sizes

Storage sizes of 512K, 1024K, 1536K, 2048K, and 3072K are provided. The Model 65 offers a maximum of 1024K. Larger Model 165 storage sizes are available at smaller cost increments, and additional storage can contribute significantly to system performance and capabilities.

The addition of more processor storage provides the Model 165 user with the ability to:

- Execute more jobs concurrently, including new application and integrated 7000-series emulator jobs
- Add and expand applications, such as graphics, teleprocessing, time sharing, remote job entry, and data based, that require large amounts of storage
- Use higher level language translators and linkage editors that provide more functions and execute faster
- Execute larger processing programs without the necessity of overlay structures
- Allocate more processor storage to language translators and sorts to improve their execution speed
- Use more and larger I/O buffers to speed up input/output operations and optimize use of direct access storage space
- Include system generation options that improve control program performance and support additional functions

Expanded Channel Capabilities

The fast internal performance of the Model 165, together with expanded use of multiprogramming, requires that more data be available faster.

The Model 165 offers more and faster channels than the Model 65, twelve instead of seven, and 1.5 MB and 3.0 MB data rates on the 2880 in addition to a 1.3 MB rate on the 2860.

The channel features of the Model 165 provide:

- A significantly higher attainable aggregate data rate than provided by the Model 65 to balance the high performance capabilities of the Model 165 CPU. A twelve-channel Model 165 system can support I/O configurations with an aggregate data rate in excess of 9 MB.
- Attachment of high-speed direct access devices, such as 3330 and 2305 facilities, that will increase I/O throughput
- Channel throughput increase capabilities via use of block multiplexing with rotational position sensing to improve effective data transfer rates

Faster I/O Devices with Increased Data Capacity

The 3330 and the 2305 facilities offer significantly faster data access than the 2314 facility and 2301 Drum Storage because of higher data transfer rates, faster rotation, and new features. Rotational position sensing and multiple requesting used with block multiplexing

can improve I/O throughput by making more efficient use of channel time. These direct access facilities also offer higher availability through use of advanced hardware-only and program-assisted error correction features.

The 3330 facility provides high capacity and fast access for less cost per bit. It is a growth device for the 2314 facility and the 2321 Data Cell Drive that offers increased price performance. The 3330 facility is designed to be used in every area in which direct access storage is needed, for example:

- As a system residence device and for program library storage
- In teleprocessing applications for message queuing and residence of online applications data
- In online, data-based applications, such as management information systems, airline reservations, etc.
- In time-sharing (or interactive) environments as online work storage (for program and data residence)
- As high-speed work storage for sorting, assembling, and link editing
- For residence of data indexes, such as for ISAM data sets

The 2305 facilities offer larger capacity and faster access than the 2301 drum. For Model 165 users, the 2305 facilities can contribute significantly to system throughput improvements when used:

- As system residence devices
- In time-sharing environments, for program and data residence
- As high-speed work storage and for residence of data indexes

SUMMARY

The combination of new and improved hardware and input/output facilities, enhanced operating systems support, integrated emulation, and increased system availability provided by the Model 165 offers Model 65 and 75 users expanded computing capabilities without the necessity of a large conversion effort. Little or no time need be spent modifying operational System/360 code or 7000-series programs currently being emulated. Existing CPU-bound programs can execute faster because of the increased internal performance of the Model 165, while I/O-bound programs can benefit from the use of more storage, more channels, faster I/O devices, and block multiplexing.

The increased power and new functions of the Model 165 provide the base for expanded applications growth and penetration of previously marginal application areas. The increased price performance of the Model 165 offers the user the opportunity to widen his data processing base for comparatively less cost and the Model 165 can be an integral part of a growth plan to the higher performance Model 195.

SECTION 10: ARCHITECTURE AND SYSTEM COMPONENTS

10:05 ARCHITECTURE DESIGN

The basic design objectives embodied in System/370 Model 165 architecture provide System/360 Model 65 and 75 users and 7000-series emulator users with a growth system that incorporates improvements and additions to System/360 architecture. The Model 165 provides new system capabilities, performance improvements, and features to increase system reliability, availability, and serviceability. This has been achieved under the following conditions:

- System/370 Model 165 architecture is upward compatible with that of System/360 models so that most user programs written for System/360 will run efficiently on the Model 165 without modification.
- Programming systems support for the Model 165 is based on that provided for System/360 models, namely on OS MFT and MVT.
- Most currently announced System/360 I/O devices will operate on the Model 165. (See Section 20:05 for a list of the I/O devices that cannot be included in a Model 165 configuration.)
- The open-ended design characteristic of System/360 has been preserved and extended in System/370.

As a result of the architecture design criteria used for this new system, all programs written for System/360 (Models 25 and up) will operate on a System/370 Model 165 with a comparable hardware configuration, with the following exceptions:

1. Time-dependent programs
2. Programs using machine-dependent data such as that which is logged in the machine-dependent logout area. (OS SER and MCH and DOS MCRR error-logging routines for System/360 models will not execute correctly.)
3. Programs that use the ASCII mode bit in the PSW
4. Programs that depend on the nonusable lower processor storage area being smaller than 1512 bytes. This area can be reduced to 512 bytes by moving the CPU logout area. (See Section 50.)
5. Programs deliberately written to cause certain program checks
6. Programs that depend on devices or architecture not implemented in the Model 165, for example, the native file of the Model 44, relocation implemented in the Model 67, etc.
7. Programs that use model-dependent operations of the System/370 Model 165 that are not necessarily compatible with the same operations on System/360 models

Note that these are the same types of restrictions that exist for compatibility among System/360 models.

The major elements of the Model 165 computing system are illustrated in Figure 10.05.1. Each component and its new features are discussed

Architecture Implementation Alterations

Two alterations have been made to the system action taken on a Model 165 during the execution of certain instructions common to both System/370 and System/360 models. The first involves all instructions that check the validity of operands involved in packed decimal operations. On the Model 165, an invalid sign in an operand causes the instruction to be suppressed (never executed) rather than terminated during execution as is done on System/360 models. Suppression, rather than termination, of an instruction when an invalid sign occurs insures that the data fields involved remain unchanged. Therefore, when a program check occurs, a routine can be executed that inspects the field that has the invalid sign.

For example, when an invalid sign results from packing an entirely blank field, the sign can be corrected by programming, and transaction deletion or program termination is avoided.

The second alteration concerns the recognition of a storage protection exception during the execution of an EDIT or an EDIT AND MARK instruction. On a Model 165 a protection exception always occurs when a pattern character is fetched from a location protected for storing but remains unchanged during the edit operation. This change eliminates unpredictable system operation during editing on a Model 165. The occurrence of a protection exception for the situation described is model-dependent for System/360 models.

Interval Timer (Standard)

The interval timer in decimal location 80 in fixed processor storage of a Model 165 has a resolution of 3.33 ms instead of the 16.6 ms resolution (with 60-cycle power) implemented for the standard timer on the Model 65. Its maximum time period remains 15.5 hours. The higher resolution of this interval timer will eliminate many of the problems encountered in accounting routine accuracy caused by task execution durations that are less than the 16.6 ms resolution.

Time of Day Clock (Standard)

This new clock is a binary counter of 52 bits with a cycle of approximately 135 years. It is updated every microsecond. Two new instructions (SET CLOCK and STORE CLOCK) are provided to set the time and to request that the current time be stored in a specified double-word of processor storage. The time can be set only when the CPU is in supervisor state and only when the clock security switch on the system console panel is in the enable set position.

The time of day clock can be used for more accurate time stamping than the interval timer. Accurate time of day can be maintained because the clock stops only when CPU power is turned off. The interval timer cannot be as accurate as the clock for time of day maintenance because it is not updated when the system is in the stopped state, and its updating may be omitted under certain conditions of excessive system activity. The 15.5-hour cycle time of the interval timer is also a restriction. The time of day clock better answers the timing needs of teleprocessing and real-time applications.

CPU Cooling

The heat generated by the logic boards in the Model 165 CPU and its associated power frames is removed by forced air and a closed-loop water circulation system. Use of a liquid coolant in addition to air is required because of the amount of heat generated by the densely packed circuits in the CPU.

The user must supply 30 gallons of cooled water per minute (45° to 60° F.) to the coolant distribution unit (CDU), which is housed in a stand-alone frame that also contains power and the power distribution unit. Water is supplied to the CDU in pipes under the raised floor. The chilled water entering the CDU is used to control the temperature of the internal water that passes through the CPU. That is, the user-supplied water does not enter the closed-loop system of the CPU. (See Figure 10.10.1.) The CDU houses the necessary controls to maintain the proper temperature in the closed-loop system. The user must supply the controls to maintain the temperature of the chilled water supplied to the CDU.

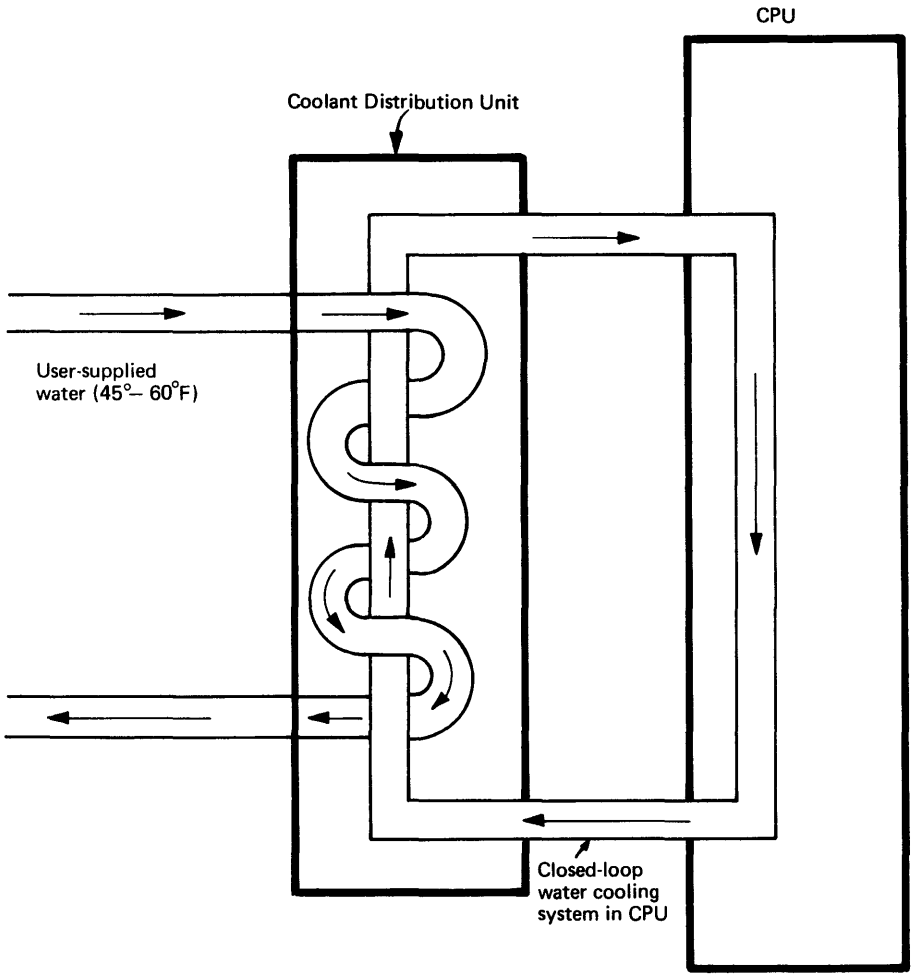

Figure 10.10.1. Conceptual flow of the water cooling system in the Model 165

The use of water as the cooling liquid offers several advantages. First, it is readily available. The chilled water normally supplied for air conditioning is acceptable. Second, water offers safety features. Low pressure is required and the cooling system can operate at room temperature, thereby eliminating problems with condensation. Last, a simplified circulation system suffices, with relatively few moving parts and less exposure to leaks. Pipe components within the CPU are separated easily, and the valved connections close off automatically to prevent water from escaping.

to the CPU without a processor storage reference. If the requested data is not currently in the buffer, a processor storage fetch is made. The data obtained is sent to the CPU. The data is also assigned a buffer location and stored in the buffer. When data is stored by the CPU, both the buffer and processor storage are updated if the processor storage location being altered is one whose contents are currently being maintained in the buffer.

The channels never access the buffer directly. They read into and write from processor storage only. When a channel stores data in processor storage, the address array is interrogated. If data from the affected processor storage address is being maintained in the buffer, the data is placed in the buffer as well as in processor storage.

The entire buffer can be disabled manually by a system console switch or via execution of a DIAGNOSE instruction. When the buffer is disabled, all CPU fetches are made directly to processor storage and effective system execution speed is reduced.

The 8K buffer is shown in Figure 10.15.4. It contains 64 columns of 128 bytes each. Every buffer column is subdivided into four blocks. A block is 32 bytes and can contain 32 consecutive bytes from processor storage that are on a 32-byte boundary. The 8K buffer can contain a maximum of 256 different blocks of processor storage data (4 blocks per column times 64 columns). A valid trigger is associated with each buffer block and is set to indicate whether or not the block contains valid data.

Processor storage is logically divided into the same number of columns as buffer storage: 64 or 128. While there are always four blocks in a buffer column, the number of blocks in a processor storage column varies with the size of processor storage. For example, when an 8K buffer is present, bits 21-26 of the processor storage address determine which of the 64 columns to use. As shown in Figure 10.15.4, a processor storage column consists of 512 blocks in a system with 1024K. Any of the 512 blocks in a given processor storage column can be placed in any of the four blocks in the corresponding buffer column. Figures 10.15.5 and 10.15.6 show the formats used for processor and buffer storage addressing.

The larger the processor storage size, the greater the number of storage blocks that contend for the four blocks in the same buffer column. If a 16K buffer is used instead of an 8K, the number of buffer columns is doubled from 64 to 128; the number of processor storage blocks contending for the blocks in each buffer column is thereby halved if processor storage size remains the same. (A 1024K processor storage divided into 128 columns has 256 blocks per column.)

The 16K buffer is provided for users with larger Model 165 processor storage configurations that have applications such that increased system throughput results from an increase in internal performance.

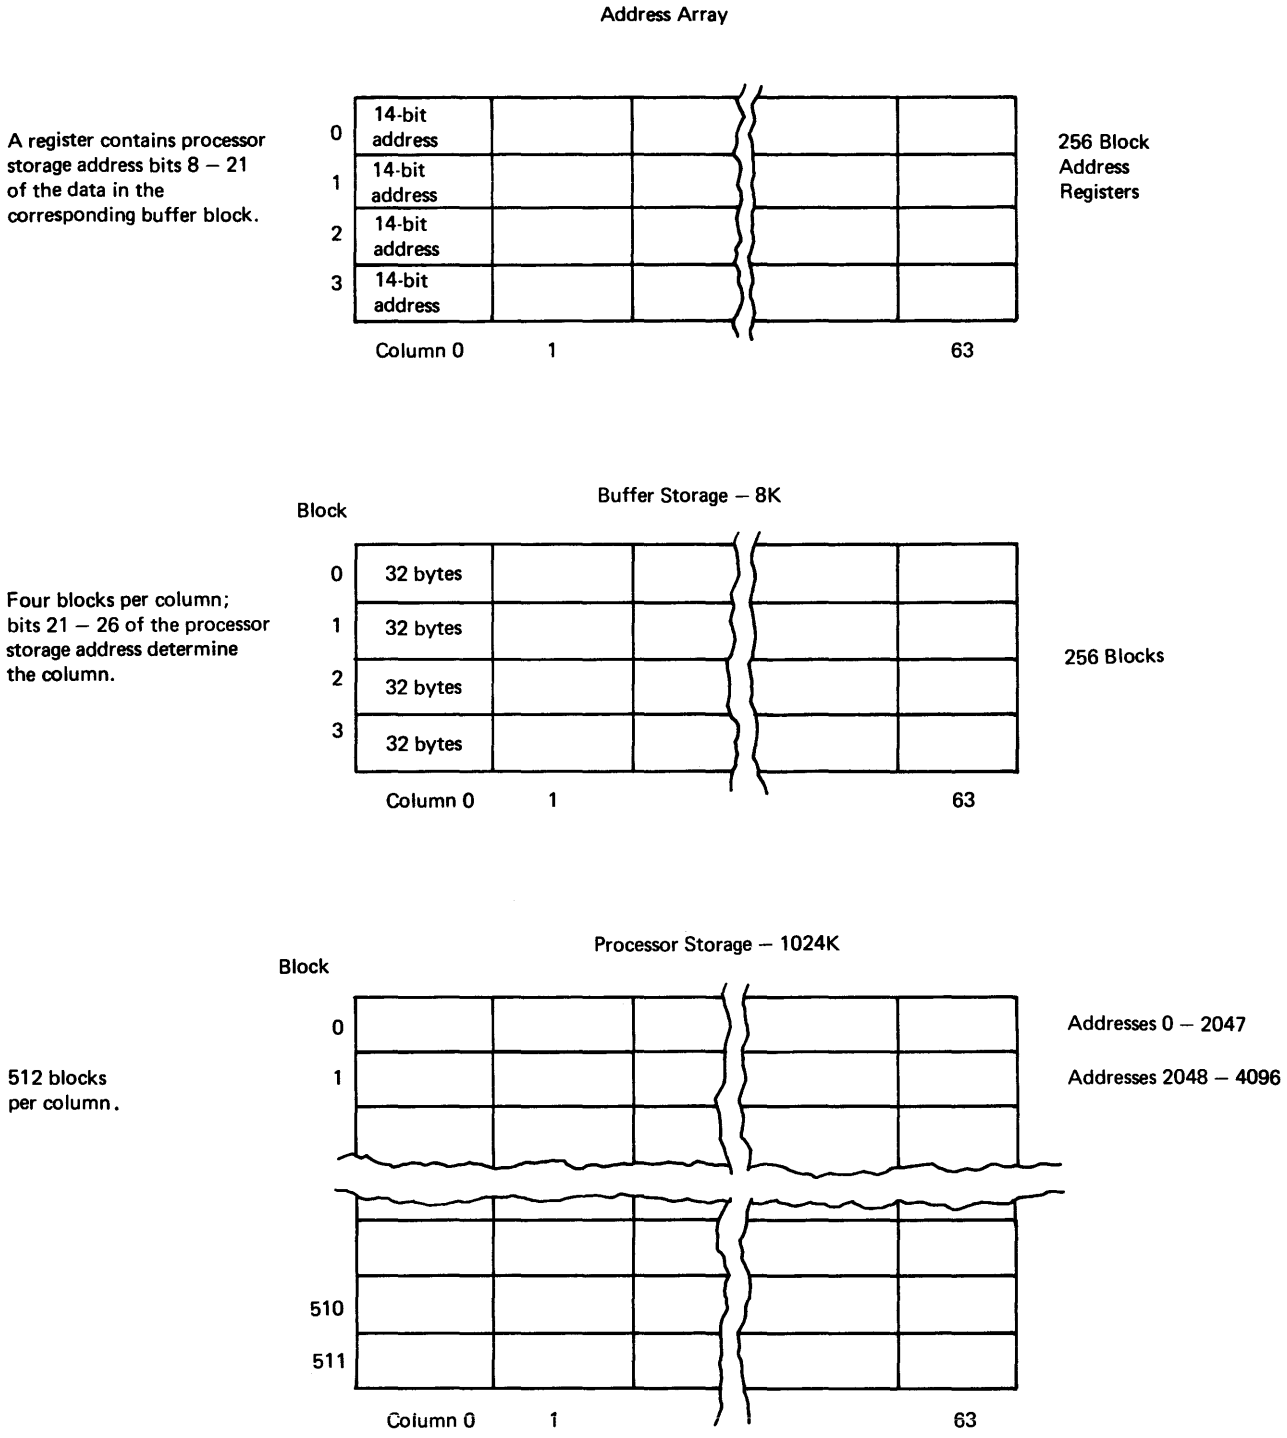

Figure 10.15.4. 8K buffer organization

Processor Storage Address Bits

8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Bits

8-21	Used for address compare
20-26	Used to reference 16K buffer columns
21-26	Used to reference 8K buffer columns
27-28	Used to reference doubleword within a block
29-31	Used to reference byte within doubleword

Figure 10.15.5. Processor storage address format for buffer reference

Buffer Reference Bits

0	1	2	3	4	5	6	7	8	9	10

Bits

0-1	Generated as result of address compare; this two-bit encoded field represents block in buffer column containing desired data
2-8	Used to reference column for 16K buffer (processor storage address bits 20-26)
3-8	Used to reference column for 8K buffer (processor storage address bits 21-26)
8-10	Used to reference doubleword within a block (processor storage address bits 27-28)

Figure 10.15.6. Buffer address format

Buffer contents and buffer block assignment are controlled by an address array, shown in Figure 10.15.4, and a special replacement array. The address array, like the 8K buffer, is divided into 64 columns consisting of four block address registers each such that there is a one-for-one correspondence between address array registers and blocks in the buffer. An address array block register contains the 14-bit processor storage block address from bits 8-21 of the processor storage address of the data contained in its corresponding buffer block. When a CPU to processor storage reference is made, the four appropriate address array column registers (14-bit block addresses) are interrogated to determine whether the requested data is currently in the buffer.

The replacement array is used to maintain knowledge of the activity of the data blocks within each of the buffer columns. When an 8K buffer is present, the array consists of 64 logic-controlled activity lists, one list for each column in the buffer. A list contains four entries, one for each buffer block in its column. A block's entry is placed at the top of the list for its column when the buffer block is referenced during a CPU fetch operation. This approach insures that the block used longest ago within a given column is at the bottom of the list. When a block within a buffer column has to be assigned and loaded, because the data requested by the CPU is not in the buffer, the buffer block at the bottom of a column activity list is allocated. Thus, the more active data is maintained in the buffer.

The 8K buffer operates as follows. When the CPU requests data, bits 21-26 of the data's processor storage address are used to obtain a buffer column address. The 14 high-order bits of the processor storage address are then compared to the address in each of the four block address registers in that buffer column in the address array.

If an equal compare occurs for an address in one of the registers in the address array and the valid trigger for that block is on in the buffer, the appropriate doubleword from the buffer block is sent to the CPU as determined by bits 27 and 28 of the processor storage address. A processor storage reference is not made. The referenced buffer block is put at the top of its column activity list in the replacement array.

If the desired data block is not in the buffer column interrogated, the requested data must be fetched from processor storage, sent to the CPU, and stored in the buffer. The replacement array entry for the column involved is inspected and the buffer block at the bottom of the activity list is assigned to receive the requested data from processor storage. Four processor storage references, one cycle apart, are made to obtain the 32 consecutive bytes and place them in the assigned buffer block. The valid trigger for the buffer block is set on and the 14 high-order processor storage address bits are placed in the appropriate column within the address array. The first doubleword fetched from processor storage is the one containing the data required by the CPU. It is sent to the CPU as well as to the buffer so that processing can continue as soon as possible.

Assuming no channel interference, a 32-byte buffer block can be filled from processor storage in 1.44 microseconds (or 18 cycles) with four-way interleaving utilized. One buffer block can be loaded every 2 microseconds (25 cycles), assuming no interference.

10:20 CHANNELS

GENERAL DESCRIPTION

The channels available on a System/370 Model 165 are functionally compatible with those of System/360 models. Combinations of 2870 Multiplexer, 2860 Selector, and 2880 Block Multiplexer channels can be attached to a Model 165. If the Extended Channels special feature is installed, up to twelve channels can be connected to a single Model 165 to provide an aggregate channel data rate in excess of nine megabytes, twice the rate possible on a Model 65.

The 2870 and 2860 channels that attach to the Model 165 are the same as those used on a Model 65 but have minor hardware changes (as is true when these channels are attached to System/360 Models 75 and up). When a Model 165 replaces a Model 65 or 75, the 2870/Model 165 and 2860/Model 165 features can be field installed on the existing 2870 and 2860 channels, respectively, so that the latter can be attached to the Model 165.

The standard number of addressable channels permitted on a Model 165 is seven. Any combination of one or two 2870 Multiplexer, up to six 2860 Selector, and up to six 2880 Block Multiplexer channels can be attached up to the limit of seven channels. If the optional Extended Channels feature is installed, the maximum number of each type of channel that can be included in the twelve channels permitted is:

- 2870 Multiplexer - two. (The first must have address 0, the second can be assigned an address from 1 to 6.)
- 2860 Selector - six. (Addresses 1 through 6 can be assigned.)
- 2880 Block Multiplexer - eleven. (Addresses 1 through 11 can be assigned.)

Model 165 channels are not integrated with the CPU. Channels compete with each other and the CPU only for processor storage accesses and, therefore, cause a minimum of CPU interference. A 2870, 2860, or 2880

channel contains the hardware required to control its I/O operations (channel registers, local storage, control functions, buffers, etc.). A channel interferes with the Model 165 CPU if the CPU accesses a logical memory that is busy because of a channel operation. Contention between channels and the CPU for processor storage is reduced drastically by the use of high-speed buffer storage, which eliminates CPU to processor storage fetches for approximately 95% of the fetches required.

The standard instruction set also includes a new I/O instruction called HALT DEVICE. This instruction is specifically designed to stop an I/O operation on a particular device on a multiplexer channel without interfering with other I/O operations in progress on the channel. HALT DEVICE should always be used instead of HALT I/O to stop an I/O operation on a 2880 Block Multiplexer channel.

The 2870 Multiplexer can control concurrent execution of 192 slow-speed to medium-speed devices, one with each of its 192 subchannels. Depending on the channel priority assigned, the 2870 can support a data rate of up to 110 KB. The maximum aggregate byte data rate of the 2870 Multiplexer subchannels is reduced by the inclusion of one or more selector subchannels, each of which can have up to 16 I/O devices attached. Each of the first three selector subchannels included can operate at a 180 KB rate, can handle one burst operation at a time, and reduces the aggregate byte rate of the multiplexer interface by 10 KB to 25 KB depending on the priority of the 2870 and the total I/O configuration. The fourth selector subchannel can operate at a rate of 100 KB and further reduces the maximum aggregate rate of the multiplexer subchannels by 14 KB. If two 2870 Multiplexer Channels are installed, the second 2870 can have only two selector subchannels.

The 2860 and 2880 channels permit attachment of a wide variety of high-speed I/O devices to the Model 165. The 2860 Selector Channel handles data rates of up to 1.3 megabytes, while the 2880 Block Multiplexer can handle a 1.5-megabyte rate, which permits attachment of the 2305 Model 2 facility. The Two-Byte Interface optional feature can be installed on a 2880 to double its data transfer rate capability to 3.0 megabytes so that the 2305 Model 1 facility can be attached.

The Model 165 user has more flexibility than the Model 65 user when configuring channels. On a Model 165, channel address and channel priority are not related, as they are on the Model 65. Any channel address (0 through 11) can have any channel priority (1-12) assigned. Thus, the Model 165 user can assign channel priority by channel type in order to achieve the desired aggregate channel rate and insure that the highest speed devices are assigned the highest channel priorities.

The maximum speed of an I/O configuration possible on a Model 165 with minimal overrun exposure is a function of the I/O devices used, the channel types installed, the channel priority assignments, and the types of channel programs operating concurrently at any particular instant. Examples of configurations that will operate on the Model 165 follow.

1. A seven-channel system including one 2870 Multiplexer and six 2860 Selector Channels with a maximum aggregate channel data rate of approximately 4.25 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2860	2301 Drum	1.250
2	2	2860	2301 Drum	1.250
3	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 3 SSC3	.090
			2540, 1403, terminals on MPX	.020
				.470
4	3	2860	2314 or 2420 Model 7	.320
5	4	2860	2314 or 2420 Model 7	.320
6	5	2860	2314 or 2420 Model 7	.320
7	6	2860	2314 or 2420 Model 7	.320

2. A seven-channel system including one 2870 Multiplexer and six 2880 Block Multiplexer channels with a maximum aggregate data rate of approximately 7.8 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2880	2305 Model 1	3.0
2	2	2880	2305 Model 2	1.5
3	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 3 SSC3	.090
			2401 Model 3 SSC4	.090
			Unit record on MPX	.023
				.560
4	3	2880	3330 facility	.806
5	4	2880	3330 facility	.806
6	5	2880	3330 facility	.806
7	6	2880	2314 or 2420 Model 7	.320

3. A seven-channel system including two 2870 Multiplexer and five 2880 Block Multiplexer channels with maximum aggregate data rate of approximately 6.8 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	1	2880	2305 Model 1	3.0
2	0	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			2401 Model 6 SSC3	.180
			2401 Model 3 SSC4	.090
			MPX	.027
				.660
3	2	2870	2401 Model 6 SSC1	.180
			2401 Model 6 SSC2	.180
			MPX	.040
				.400
4	3	2880	3330 facility	.806
5	4	2880	3330 facility	.806
6	5	2880	3330 facility	.806
7	6	2880	2314 or 2420 Model 7	.320

4. A twelve-channel system including one 2870 Multiplexer and eleven 2880 Block Multiplexer channels with a maximum aggregate data rate of approximately 9.4 MB.

Channel Priority	Channel Address	Channel Type	I/O Device	Data Rate in MB
1	8	2880	2305 Model 1	3.0
2	9	2880	2305 Model 2	1.5
3	0	2870	Multiplexer aggregate	.560
4	10	2880	3330 facility	.806
5	11	2880	3330 facility	.806
6	1	2880	3330 facility	.806
7	5	2880	2314 or 2420 Model 7	.320
8	3	2880	2314 or 2420 Model 7	.320
9	4	2880	2314 or 2420 Model 7	.320
10	2	2880	2314 or 2420 Model 7	.320
11	6	2880	2314 or 2420 Model 7	.320
12	7	2880	2314 or 2420 Model 7	.320

The Model 165 with a 2 microsecond processor storage is able to achieve higher aggregate data rates than the Model 65 with a .750 microsecond main storage because Model 165 operations have been designed to optimize the use of four-way doubleword interleaving by use of more I/O buffering. Additional channel buffers are provided in the Model 165 CPU and the 2880 channel, which permit effective use of interleaving during I/O operations and help prevent data overrun when peak activity occurs.

The storage control unit contains four dedicated 8-byte buffers for each channel that are assigned when the channel is installed. In addition, each 2860 Selector Channel contains two 8-byte buffers, and a 2880 Block Multiplexer*Channel contains two 16-byte buffers or four 16-byte buffers depending on its maximum rate, 1.5 MB or 3 MB.

The use of I/O buffers in the SCU allows channel requests to be made to processor storage concurrently with other channel or CPU requests when requests are for different logical memories. Additional internal buffering in the 2880 channel allows it to withstand longer wait times for requested logical memories and thus reduces channel overrun exposure.

The channel buffering scheme implemented in the Model 165 is most efficient for 2880 channels. Use is made of all four buffers available per channel, while only two buffers can be used for 2860 channels. This scheme supports overlapped processor storage requests from the same 2880 channel so that at any given channel priority position, a 2880 can sustain a higher data rate than a 2860.

A 2301 drum connected to a 2860 Selector Channel must be placed in the highest or second highest channel priority position on a Model 165. Channel address 1 or 2 must be used. For most efficient operation of 2301 drums, they should be attached to a Model 165 via a 2880 channel, since the 2301 causes much less CPU interference during data transfer operations when connected to a 2880, rather than to a 2860.

In contrast to the Model 165, the Model 65 does not contain additional buffering in its CPU and because of timing considerations cannot utilize the two-way interleaving capability of its main storage for I/O operations. Thus, in a Model 65, without the use of buffering in the SCU, the memory bus is busy for an average of 1.1 microseconds when storing one doubleword from a 2860 channel buffer. In a Model 165, with no interference, four doublewords from channel buffers in the SCU can be placed in processor storage in 2 microseconds (one doubleword in each logical memory).

Comprehensive error checking has been incorporated in the basic design of the channel hardware. Checking is performed on the control logic in most areas and standard parity checking is done on the data flow. Improved error recovery hardware has also been included (discussed fully in the RAS section).

The Channel-to-Channel Adapter feature available for the Model 65, which permits two System/360 channels to be interconnected, is also an optional feature for the Model 165. The Channel-to-Channel Adapter itself cannot be installed on a 2880 or a 2870 channel; however, a 2880 or a 2870 can be connected to an adapter installed on another channel. Thus, a 2880 or a 2870 can be interconnected to a 2860 channel, a Model 155 channel, a Model 50 channel, etc., that has the Channel-to-Channel Adapter attached.

THE 2880 BLOCK MULTIPLEXER CHANNEL

The 2880 can operate either as a selector or a block multiplexer channel. The setting of a channel mode bit in a control register determines the mode in which 2880 channels operate. The mode bit is set to selector mode at IPL and on system reset and can be altered by programming at any time. When a START I/O instruction is issued to a 2880, the setting of the channel mode bit determines the mode in which the addressed subchannel will operate.

The new START I/O FAST RELEASE instruction can be used with 2880 Block Multiplexer channels. This instruction differs from a START I/O in that START I/O FAST RELEASE permits the CPU to execute the next sequential instruction sooner. That is, if the addressed 2880 is not busy when START I/O FAST RELEASE is issued, the assumption is made that the I/O operation can be started and the CPU is not held up awaiting a response from the device's control unit. If it is determined

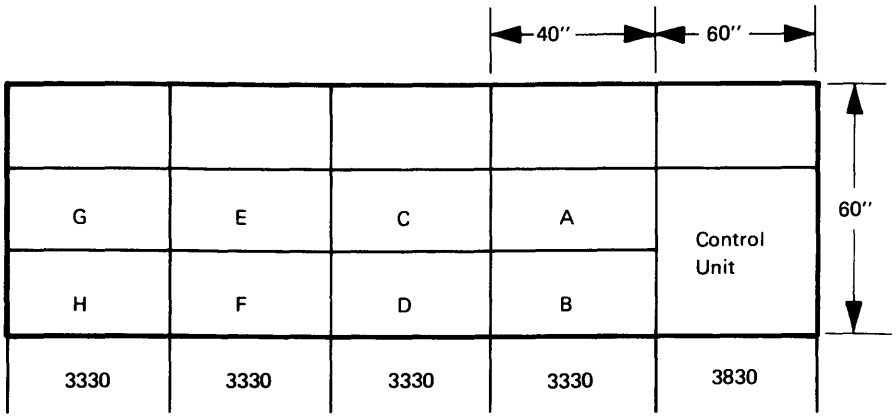

Figure 20.10.1. The 3330 facility

Facility configurations and maximum capacities, using full track records, are shown below:

3830 Storage Control + one 3330 module	200 megabytes
3830 Storage Control + two 3330 modules	400 megabytes
3830 Storage Control + three 3330 modules	600 megabytes
3830 Storage Control + four 3330 modules	800 megabytes

Drives are mounted in powered drawers that are opened and closed by a switch on the operator control panel on the 3330 module. Logical address plugs are supplied, as for the 2314, in addition to a CE service plug. The latter is used when inline diagnostics are to be executed.

Functionally, the 3330 facility provides more capabilities than the 2314, particularly in the areas of performance and availability. The 3330 facility supports all the standard 2314 commands (except the file scan commands) in addition to several new operations, including RPS and error recovery commands. The 3330 facility also is an attractive growth device for the 2321 Data Cell Drive.

The new, removable 3336 Disk Packs are interchangeable across 3330 disk drives but are not interchangeable with the 2316 Disk Packs used on 2314 disk drives (Table 20.10.2 compares disk pack characteristics). Like 2316 packs, 3336 Disk Packs will be initialized in the factory with home addresses and capacity records (R0). Up to 20 defective tracks per pack will be flagged and have alternates assigned. The QUICK DASDI utility, currently available for processing 1316 and 2316 packs, will support 3336 packs. QUICK DASDI writes the volume label, the VTOC, and, if requested, IPL records but bypasses track analysis. It also determines the number of flagged tracks and places this data in the VTOC.

Table 20.10.1 compares the capacity and timing characteristics of the 3330 facility with those of the 2314 facility and the 2321 Data Cell Drive. The increase in capacity achieved by replacing a 2314 or a 2321 with a 3330 depends upon the block size chosen for the data on the 3330. For example, if the 2314 full track block size of 7294 bytes is maintained for a given data set on the 3330 to avoid programming changes, the 3330 yields a 91% increase in full pack capacity (almost twice the capacity). However, reblocking to a full track on the 3330, 13,030 bytes, yields a 242% full pack capacity increase. If there is not enough processor storage available to allocate I/O areas of 13,030 bytes, lowering the 3330 block size used to half a 3330 track yields a 239% increase in full pack capacity.

If a 2321 is replaced by a 3330, six full track blocks of data from the 2321 (2000 bytes/2321 track) can be placed on each 3330 track if full track blocking is used, for a total of 92,112,000 bytes per 3336 pack (12,000 bytes x 7676 tracks per 3336). Thus, slightly over four 3336 packs provide the capacity equivalent of ten data cells, or a full 2321 drive, if full tracks are used. A full ten data cells, blocked full track, also can be contained in slightly more than four 3336 packs if half-track blocking is used on the 3336.

Table 20.10.1. Capacity and timing characteristics of the 3330 and 2314 facilities and the 2321 Data Cell Drive

Characteristics	3330	2314	2321
Capacity in bytes truncated to the nearest thousand (full track records)			
Pack or cell	100,018,000	29,176,000	39,200,000
Facility or Data Cell Drive			
2 drives/cells	200,036,000	58,352,000	78,400,000
4 drives/cells	400,073,000	116,704,000	156,800,000
6 drives/cells	600,109,000	175,056,000	235,200,000
8 drives/cells	800,146,000	233,408,000	313,600,000
10 cells	-	-	392,000,000
Access time in ms			
Maximum	55	130	600 (for strip select and load)
Average	30	60	175 (minimum for strip select and load)
Average cylinder to cylinder	10	25	95 (on a strip)
Rotation time in ms	16.7	25	50 (strip on drum)
Rotation speed (rpm)	3600	2400	1200
Data transfer rate (KB)	806	312	55

Table 20.10.2. 3336 and 2316 Disk Pack characteristics

Characteristics	3336	2316
Number of disks per pack	12	13
Number of recording disks	10	11
Number of recording surfaces (recorded tracks per pack)	19	20
Disk thickness in inches	.075	.050
Disk diameter in inches	14	14
Disk pack weight in pounds	20	15
Disk pack maximum capacity in millions of bytes	100	29.1
Full track capacity in bytes	13,030	7294
Cylinders per pack	404 plus 7 alternates	200 plus 3 alternates
Tracks per cylinder	19	20
Tracks per pack	7676	4000

Self-formatting records are written on 3336 packs in the same manner as they are on 2316 packs. However, each physical area written (count, key, and data) has a field of error correction code appended to it for the purpose of data validity checking by the control unit instead of the cyclic check area used on the 2314.

mechanism is self-adjusting and automatically rises in increments as the stack of paper mounts. This insures that the stacker mechanism is always the same distance above the top of the stack of forms. The rate of rise during each increment is determined by the setting of the stacker rate knob, which can be adjusted by the operator to produce the best condition for the thickness of the forms being stacked. The stacker also can be raised or lowered manually.

When forms are inserted, the printer platen automatically positions itself close to the train cartridge in accordance with the thickness of the forms. Thus, correct clearance between the platen and the cartridge is achieved without operator intervention. Because of its automatic forms thickness sensing, the 3211 is sensitive to forms with a different degree of thickness at each edge. (For forms limitations, see Form-Design Considerations-System Printers, GA24-3488.)

Forms control paper carriage tape loading and unloading by the operator is eliminated by implementation of a tapeless carriage feature for the 3211. Forms spacing and skipping are controlled by a program-loaded forms control buffer (FCB) contained in the 3811 control unit.

The FCB contains 180 storage positions, each of which corresponds to a print line, that is, a single space of the carriage. Thus, forms up to 22.5 inches in length can be accommodated at 8 lines per inch spacing (or 24 inches at 6 lines per inch). Up to twelve channel codes (1-12), corresponding to the twelve channel positions of the paper carriage tape used on a 1403 Printer, can be stored in the appropriate buffer line positions to control carriage skipping. The FCB can be considered to contain a storage image of a carriage control tape.

A carriage control address register is used to address the FCB and maintain correct line position with respect to the form. This register is incremented as space and skip commands are issued that cause the form to advance. When a SKIP TO CHANNEL command is executed, the carriage control address register is incremented and the form moves until the channel specified is sensed in a line position in buffer storage. If the requested channel number is not found in the buffer, forms movement stops after address position 1 (line 1) has been sensed twice. This prevents runaway forms skipping.

A flag in a buffer storage line position is used to indicate the last line of the form for forms shorter than 180 lines. A flag bit is also used in the first buffer storage position to indicate six or eight lines per inch spacing. The FCB is loaded with the desired forms spacing characters via a LOAD FCB command issued by a program. An error indication is given if an end-of-page flag is not present or if an invalid carriage code is loaded.

Serviceability features, in addition to those provided for the 1403 Printer, are incorporated into the design of the 3211. The fact that a 3811 control unit controls only one 3211 Printer, instead of multiple devices, permits offline repair of the malfunctioning printer or control unit only, without the necessity of removing other operational units from the system.

The 3811 control unit presents six bytes of sense information to identify printer and control unit malfunctions instead of only one, as is provided for the 1403. Certain errors (such as a parity check in the print line buffer) that might be corrected by programmed retry of the print operation are identified in the sense bytes, and carriage motion is suppressed. This permits error recovery without operator intervention if the retry is successful. The additional status data presented can be stored for later analysis and should help speed the diagnosis of hardware malfunctions.

SECTION 30: PROGRAMMING SYSTEMS SUPPORT

30:05 TRENDS IN DATA PROCESSING AND PROGRAMMING SYSTEMS

The Model 165 and its programming systems support have been designed to operate in the data processing environment that has been emerging since introduction of System/360.

Significant trends are the following:

- Growth toward more multiprogramming to improve system throughput. Multiprogramming also permits the user to install new applications, such as small teleprocessing inquiry, graphics, or time-sharing applications, that would not otherwise justify a dedicated system. Multiprogramming support also has encouraged the growth of new computer environments, as indicated by the items that follow.

- Growth of integrated emulation, that is, concurrent native and emulation mode processing on one system. The execution of emulators under operating system control improves system throughput because emulators can use control program facilities (stacked job execution, data management functions, etc.) and because native mode and emulator jobs can be scheduled to operate concurrently to utilize available system resources efficiently. The use of integrated emulators eliminates most reprogramming and eases transition from one system to another, permitting the user to expend his efforts extending and adding applications.

- Greater use of high-level languages, such as COBOL, FORTRAN, and PL/I, for applications programming. The cost of programming has been increasing while the cost of computing hardware has been decreasing. More productive use of programmers can be achieved by the use of high-level languages. Improvements to compile times and to the size and execution speed of code produced by high-level language translators have been made and continue to be made. The support of many more functions within high-level languages also has increased their use, and the growth of interactive computing has stimulated the addition of even more facilities.

- Growth of teleprocessing applications such as remote inquiry, message switching, data collection, and management information systems. The ability of System/360 and System/370 to handle teleprocessing and batch processing in one system eliminates the necessity of dedicated special purpose systems.

- Growth of remote computing activities, such as remote job entry and interactive computing (or time-sharing) in both a nondedicated and dedicated environment. Remote computing offers (1) fast turnaround for batch work submitted from remote locations, (2) remote user access to the large computing facilities and data base available at the central installation, and (3) interactive problem solving on a regular or a nonscheduled basis for personnel in locations remote from the central computer. In-house interactive computing is growing also as users attempt to use programmer time more efficiently.

- Growth toward large, online data base systems. The growth in the marketplace of remote computing, time sharing, and real-time applications necessitates the instant availability of more and more data. High-capacity, fast, reliable direct access devices are required for this type of computing environment.

IBM programming systems support of these trends in data processing has been growing and continues to be extended. The System/370 Model 165 offers the hardware, I/O devices, and performance capability required by the expanding computing environment.

30:10 OS SUPPORT

It is expected that the majority of the users who upgrade to the Model 165 will use OS as their operating system. The design of OS and the facilities it currently provides make it a particularly suitable base upon which to build additional support of the data processing trends discussed. Some of the major OS features currently available or announced are the following:

- Priority and job class scheduling for increased resource utilization
- Multiprogramming support of up to 15 user jobs and multitasking facilities
- Automatic data transcription performed by the control program (reader interpreters and output writers) concurrently with user job execution
- Dynamic resource allocation by the control program
- Multiple console support (MCS) and device-independent display operator console support (DIDOCS)
- Extensive teleprocessing and graphics support
- Remote job entry and conversational remote job entry
- Support of a wide range of interactive (time-sharing) environments

OS will be modified and extended in future releases so that it supports Model 165 hardware. Appropriate alteration of the resident portion of a control program (nucleus) generated for a Model 165 will accommodate the fixed storage area of lower storage in the Model 165. OS for the Model 165 includes currently announced OS facilities and contains additional support to handle new Model 165 hardware features and I/O devices. Emphasis has also been placed on extending error recovery, recording, diagnostic, and repair procedures.

OS support of Model 165 features will be provided as follows (Programming systems support of RAS features is discussed in Section 50).

New instructions. The Assembler F (Type I) and Assembler H (program product) language translators will include mnemonics for the general purpose and other new instructions for the Model 165 so that these instructions can be used in user-written Assembler Language programs. The currently available OS high-level language translators will not generate the six new general purpose instructions.

Extended precision floating point. Assemblers F and H will include support of the extended precision floating-point data format and instructions. In addition, extended precision will be supported by the FORTRAN H program product.

The implementation of extended precision support by FORTRAN H is such that the language translator and the processing programs generated to include extended precision operations can operate on a System/370 or System/360 with or without extended precision hardware. The language translator contains extended precision instructions and generates them

for processing programs that use extended precision data. A program check interrupt occurs if an extended precision instruction is executed and the feature is not present in the system. This interrupt causes the processing program to call in a subroutine to handle extended precision operations. Extended precision divide is always simulated, because the extended precision facility does not include such an instruction.

Interval timer. The interval timer will be supported for the same functions as it is currently except for time of day values.

Time of day clock. This clock will be supported in MFT and MVT environments for time of day requests made by system and user tasks via the TIME macro. At IPL, the operator will have the option of validating the clock time and correcting an invalid value.

Byte boundary alignment. The programmer has the ability to byte-align binary and floating-point data in Assembler Language programs, in ANS COBOL programs (by omitting the SYNCHRONIZED clause), in PL/I programs (by specifying the UNALIGNED attribute), and in FORTRAN programs. Note that OS still expects the parameters passed to it to be properly aligned.

7080, 7070/7074, 709/7090/7094/7094II Compatibility Features. Three separate, mutually exclusive emulator programs will be provided. Emulator programs are discussed in Section 40.

Channels. One or two 2870 Multiplexer Channels and up to twelve high-speed channels (2860 and 2880) in a system will be supported.

Both selector and block multiplexer mode of operation will be supported for 2880 channels. During IPL, block multiplexer mode is established via a control register bit setting (discussed in Section 10:20) if the user requested block multiplexer mode for at least one 2880 channel at system generation. The operator does not have the option of changing this mode at IPL, nor does the control program alter the mode during system operation.

The 3330 facility. The 3330 facility will be supported as an I/O device for most of the same functions as is the 2314 facility. The error recovery routine provided for the 3330 will include support of the new hardware error correction features.

RPS will be supported in MFT and MVT environments as follows.

- The stand-alone seek issued within the I/O supervisor (IOS) will be eliminated for RPS devices (3330 facilities). IOS will continue to issue stand-alone seeks for direct access devices without RPS. IOS will also be capable of recognizing the channel available interrupt.

- Access method support of RPS commands (SET SECTOR and READ SECTOR) will be provided by:

  QSAM and BSAM for all record formats and functions except the undefined track overflow record format
  ISAM for all updating and verification and for QISAM sequential processing
  BPAM for processing directory and member processing operations where possible
  BDAM for direct retrieval and update of fixed-length standard and VBS format records without key and for write verification of all BDAM format records

- End-of-volume (EOV) routines will support concatenation of data sets residing on RPS and on non-RPS devices. The control program will ensure that an RPS access method is used for drives with the feature and that a non-RPS method is used for drives without the feature.

- Any system utility, data set utility, or IBM-supplied processing program (such as a language translator) that uses the sequential access methods will support RPS. In addition, IEHMOVE, IEBCOPY, and the initialize/dump/restore utility (IEHDASDR) will include RPS support for 3330 facilities.

- The Sort/Merge program product supports RPS for 3330 intermediate work devices.

- Where appropriate, RPS commands for access to certain SYSRES data sets will be supported by:

  Data set catalog routines
  Direct access device space management (DADSM) routines (for DSCB processing)
  STOW, BLDL, and FIND processing of program library (PDS) directory entries
  OPEN/CLOSE/EOV processing of JFCB's in the job queue
  Routines that access the job queue

  Note that RPS command support is not provided for:

  1. The telecommunications access method (TCAM) for message queues
  2. The stand-alone disk initialization and alternate track assignment routines (DASDI and ATLAS)

The 2305 facility. The 2305 facility will be supported as an I/O device for most of the same functions as is 2301 Drum Storage. Note that specification of exchange buffering for a QSAM data set on either a 2305 or 2330 facility results in a default to simple buffering.

Rotational position sensing for 2305 facilities will be supported as discussed for the 3330 facility except that an arm positioning seek is not required on 2305 modules and the 2305 is not supported as an intermediate work unit by Sort/Merge programs. Multiple requesting will be handled by the I/O supervisor, which will initiate up to seven channel programs on one 2305 module at a time. (The eighth logical device address is reserved for control purposes.) Multiple I/O operations to the same data set will be initiated if the data set is not being sequentially processed. (One operation must complete successfully before the next can be initiated during sequential processing.)

3211 Printer with tapeless carriage. The 3211 Printer, with or without the 18 additional print positions, will be supported by QSAM and BSAM for exactly the same functions as is the 1403 Printer. In addition, the control program will handle loading of the forms control buffer (FCB) with carriage control images. This support is similar to that provided for Universal Character Buffer (UCB) loading.

The user can define one or more default FCB images at system generation time. Two IBM-supplied default images are included automatically. All other FCB images to be used must be defined by the user and placed in SYS1.SVCLIB, as is the case with UCB images. User-supplied default images must be identified as defaults. The FCB image to be used by a processing program can be specified in the 3211

Printer DD statement included for the job step and will be loaded into the FCB by the control program.

The FCB image currently loaded can also be changed by the programmer during execution of the processing program by use of an Assembler Language macro. If the DD statement does not specify an FCB image and the image currently loaded is not one of the defaults specified at system generation, the operator is requested to specify the FCB image to be used.

The FCB image (and the UCS character image) to be loaded for a 3211 Printer used by an output writer can be indicated in the output writer procedure or in the START output writer command issued by the operator. FCB and UCB images can also be specified in the SYSOUT DD statement for a data set that is to be printed by the output writer, and they will be loaded into the 3811 control unit prior to the printing of the data set.

Any time the FCB parameter is used, as described above, the user can specify that operator verification of forms alignment is to be requested by the control program via a console message when the buffer is loaded. The operator must respond to this message.

The 3211 error recovery routine will retry a print operation after a parity check occurs in the UCB or print line buffer if QSAM is used and if three I/O buffers are provided for the printer data set. When the operation is retried, the 3811 control unit insures that only the print positions that did not print correctly the first time are reprinted.

Operator console. The alphameric keyboard and CRT will be supported as the primary operating system console. The multiple console support (MCS) and device-independent display operator console (DIDOCS) options of OS must be included in the operating system to support this console.

Table 40.15.2. 7094 I/O devices and features not supported by the 7094 Emulator program for the Model 165

7909 Data Channel and attached devices:
1414 Input/Output Synchronizer Model 6
7631 File Control and 1301, 1302 Disk Storage
7640 Hypertape Control and 7340 Hypertape Drives
7320 Drum Storage
7740 Communication Control System
7750 Programmed Transmission Control
Data Channel Switch
Direct data connection
Direct couple
External signal trap
740 CRT Recorder (709 device)
780 CRT Recorder (709 device)
7094 special and custom features
704 Compatibility Feature
SET DENSITY instruction (mixed density tapes)

System/370 and 7094 emulator jobs can be intermixed in the OS input stream, with special processing necessary sometimes for binary format jobs and data. OS job control and any required 7094 emulator control statements must be supplied for each 7094 program to be emulated. Existing 7094 object programs and control statements need not be modified unless they use facilities not supported by the Model 165 7094 Emulator program.

BCD format 7094 jobs and data can be placed in the input stream or in a card or tape data set. Binary format 7094 jobs in card form (object programs and data) can be placed in the input stream (in a card reader with the column binary feature) in an MVT environment if the automatic SYSIN batching (ASB) reader interpreter is used. Binary card input can also be read from a dedicated card reader with the column binary feature. Required emulator control statements can be entered via the operator's console or a data set identified by a specially named DD statement (SYSEMCTL). This data set can be in the input stream or in a card, tape, or disk sequential data set (including a member of a partitioned data set).

The 7094 Emulator program operator commands available can be entered via the operator's console or the input stream (located in the emulator control statement data set). These commands provide simulation of 7094 console operations and allow the operator to request dumps of emulator processor storage areas and to obtain certain displays (sense switch settings, register values, etc.). If multiple console support (MCS) is included in the OS control program generated, emulator program messages can be routed to a specific console device (assigned routing code 12) so that emulation messages are isolated.

Assuming that the 7094 programs to be emulated do not use facilities or I/O devices unsupported by the Model 165 7094 emulator, the following is required:

- OS job control and any required Model 165 7094 emulator control statements must be supplied for each 7094 object program to be emulated. Changes to BCD format data is not required. If binary card input is not handled by card reader operations, a utility program must be written to prepare an acceptable tape from the binary input. A user-written data set writer is required to print and punch 7094 files written to SYSOUT data sets, if any. (See Emulating the 709, 7090, 7094, 7094II on the IBM System/370 Model 165 Using OS/360, GC27-6951, for suggestions regarding these utilities.)

- Tapes containing files with records longer than 32,755 bytes (5461 7094 words) must be preprocessed by the tape formatting program. Tapes written in seven- or nine-track 7094 mode can be used without modification. (Consideration should be given to reblocking files with short blocks in order to improve execution performance and reblocking files with long blocks to reduce processor storage buffer requirements.)

- Those 7094 programs that use unsupported I/O devices or facilities can be modified to remove the unsupported facility and then emulated, or they can be rewritten as System/370 programs.

The internal performance (that is, the speed of performing 7094 CPU instructions weighted by frequency of use) of the Model 165 7094 Emulator program is approximately 1.5 times that of the IBM 7094II. The throughput of emulated 7094 jobs executed on the Model 165 will vary depending on job characteristics, resources available, etc. Total system throughput increases should be achieved by operating 7094 emulator jobs in a multiprogramming environment.

40:20 7080 EMULATOR PROGRAM

The 7080 Emulator program for the Model 165 requires a Model 165 with the 7080 Compatibility Feature (#7118), 512K or more of processor storage, and enough I/O devices for the operating system and emulated devices. In an MVT environment a 1024K minimum system is required to emulate a 7080 system with 160,000 positions of main storage. The emulator program requires a minimum processing partition or region of 240K for emulation of a 7080 system with 80K. A 364K partition or region is required to emulate a 7080 system with 160K. These two minimums support two channels and 14 tape units with one 2000-byte buffer assigned to each tape data set. (Note that other considerations, such as double buffering, I/O device configuration, performance, etc., will often necessitate use of an emulator partition or region larger than these minimums.)

Table 40.20.1 lists 7080 system features that are supported and Table 40.20.2 indicates those that are unsupported. The 7080 emulator also can execute 705 I, II, and III programs that can be executed on a 7080, subject to the considerations and limitations outlined in the IBM 7080 Principles of Operation (GA22-6560) and in this document. The Model 165 emulator supports the same facilities as the Model 65 stand-alone 7080 emulator.

The 7080 Emulator program accepts and produces two tape data formats, using BSAM:

1. 7080 format tapes, written in BCD, that are written by a 7080 system, the Model 65 stand-alone 7080 emulator, the Model 165 integrated 7080 emulator, or the Tape Postprocessor program. These tapes can contain imbedded tapemarks but cannot contain the reader storage mark character. Mixed density, seven-track input tapes can be handled if single buffering is used and the tape density specified in the DD statement is that of the first record on the tape.

   Tapes in 7080 format can be processed on nine-track tape drives and 2400 tape drives with a seven-track head attached to a control unit with the Seven-Track Compatibility Feature. They must also be designated as unlabeled as far as OS is concerned and their record format must be specified as undefined.

SECTION 50: RELIABILITY, AVAILABILITY, AND SERVICEABILITY (RAS) FEATURES

50:05 INTRODUCTION

With the growth of more and more online data processing activities, as distinguished from traditional batch accounting functions, the availability of the data processing system becomes a very essential factor in company operations, and complete system failure is extremely disruptive. Because of the growing frequency of online processing and the fact that the System/370 Model 165 is designed to operate in such an environment, IBM has provided an extensive group of advanced reliability, availability, and serviceability features for the Model 165. These RAS features are designed to improve the reliability of hardware, to increase the availability of the computing system, and to improve the serviceability of system hardware components.

The RAS features of the System/370 Model 165 are designed to reduce the frequency and impact of system interruptions that are caused by hardware failure and necessitate a re-IPL as follows.

- More reliable components, such as integrated circuits with fewer connections, will be used to improve hardware reliability.
- Recovery facilities, both hardware and programming systems, not available on System/360 Models 65 and 75, are provided to reduce considerably the number of failures that cause a complete system termination. This permits deferred maintenance.
- Repair procedures include more online diagnosis and repair of malfunctions concurrently with normal job execution in a multiprogramming environment in order to reduce the effect of such repairs on system unavailable time.

Each RAS feature, recovery or repair, is discussed in the remainder of this section.

The following recovery features are implemented in hardware:

- CPU retry of most failing CPU operations, including those caused by a buffer malfunction
- ECC validity checking on processor storage to correct all single-bit errors
- I/O operation retry facilities, including channel retry data and channel/control unit command retry procedures, to correct failing I/O operations
- Expanded machine check interrupt facilities to facilitate better error recording and recovery procedures

The following recovery features are provided by programming systems:

- Recovery management support (RMS) to handle the expanded machine check interrupt and channel retry data. Model 165 MCH and CCH routines are provided for OS MFT and MVT only.
- Error recovery procedures (ERP) to retry failing I/O device and channel operations

- OBR and SDR routines to record statistics for all temporary and permanent I/O errors
- Environment recording, edit, and print program (EREP) to format and print error log records
- I/O RMS routines, alternate path retry (APR), and dynamic device reconfiguration (DDR) to provide additional recovery procedures after channel or I/O device failures
- Advanced checkpoint/restart and warm start facilities to simplify and speed up system restart procedures after a failure necessitates a re-IPL

The following repair features are provided:

- Online Test Executive Program (OLTEP) and Online Tests (OLT's) that execute under OS control and provide online diagnosis of I/O device errors for most devices that attach to the Model 165
- Processor Logout Analysis program that operates under OS control to analyze machine check error records in order to determine suspected malfunctioning field-replaceable units
- System Test, Channel Test, CPU Test, and Storage Test stand-alone diagnostic programs to identify failing hardware units
- Microdiagnostics for customer engineer use to locate the field-replaceable unit within a malfunctioning component

These aids are designed to enhance system availability. In many cases, the system can run in a degraded mode so that maintenance can be deferred to scheduled maintenance periods. When solid failures do occur, their impact can be reduced by faster isolation and repair of the malfunction than is possible currently.

50:10 RECOVERY FEATURES

Additional hardware that attempts correction of most hardware errors without programming assistance has been included as part of the basic Model 165 system. The control program can be notified, via an interrupt, of both intermittent and solid hardware errors so that error recording and recovery procedures can take place.

AUTOMATIC CPU RETRY

Detected CPU hardware errors, except those that occur during the execution of DIAGNOSE, TEST AND SET, READ DIRECT, WRITE DIRECT, and certain instructions that have passed beyond a threshold point, can be retried automatically by CPU retry hardware. A mask bit in a control register determines whether the CPU retry function is enabled or disabled. If enabled, retry occurs after instruction errors, after failures that occur during interrupt time when status information is being saved, after errors that occur during status saving for I/O instructions, etc. An I/O instruction, such as START I/O or TEST I/O, can be retried automatically by the hardware without an intervening I/O interrupt if the instruction has not proceeded beyond an established threshold point.

CPU retry also occurs when an instruction error results from a buffer malfunction, if the instruction is a retryable type. The buffer is bypassed while the instruction is retried so that processor storage

Fixed Logout Area Locations 232 - 239

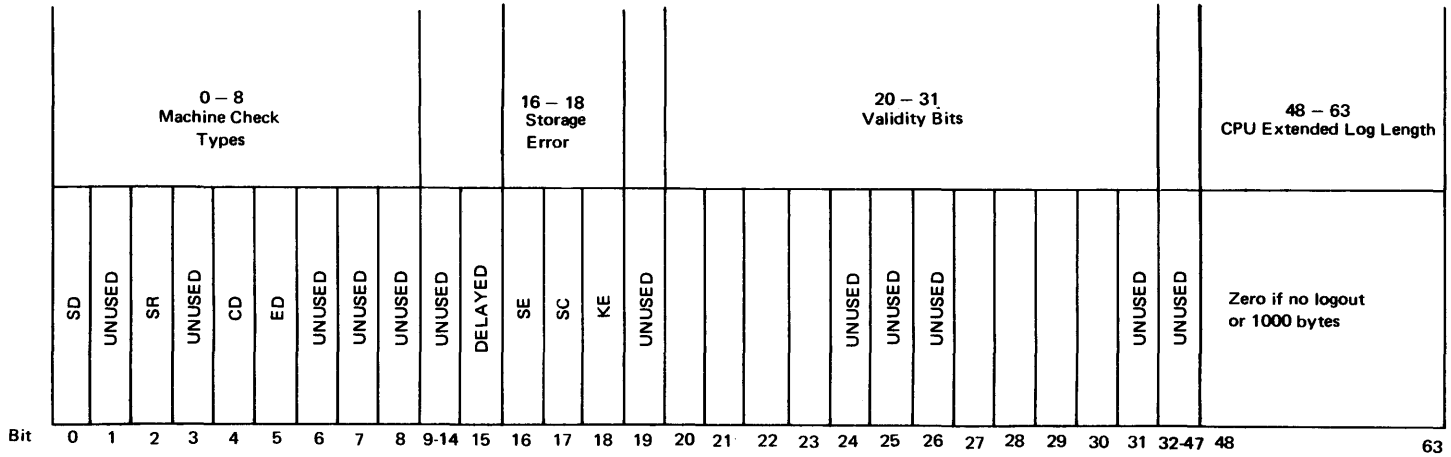

Bit	Inrerrupt Type
0	SD – System Damage
2	SR – System Recovery
4	CD – Time of Day Clock Damage
5	ED – External Damage

Bit	Error
15	Delayed Interrupt
16	Storage Error Uncorrected
17	Storage Error Corrected
18	Protection Key Error

Bit	Valid Fixed Area Data
20-23	Machine Check Old PSW (48-55)
	20 AMWP
	21 Masks and Protect Key
	22 Program Mask and Condition Code
	23 Instruction Address
27	Floating Point Registers (352-383)
28	General Purpose Registers (384-447)
29	Control Registers (448-511)
30	CPU Extended Logout

Figure 50.10.4. Machine check code - Model 165

- External Damage - This interrupt occurs if both PSW bit 13 and the external mask bit are on. The ED bit of the stored machine check code (bit 5) will be on to indicate that the storage control unit or processor storage is damaged but that operations associated with the CPU were not affected because the multiple-bit storage failure was associated with an I/O operation. When this type of error occurs, both a soft machine check and an I/O interrupt occur. The error is logged after the machine check interrupt. Error recovery procedures are initiated after the I/O interrupt.

- Time of Day Clock Damage - This interrupt occurs if both PSW bit 13 and the external mask bit are on. The CD bit in the stored machine check code (bit 4) will be on to indicate that an error occurred in the time of day clock that renders the clock invalid. Once this invalid indication has been given, subsequent STORE CLOCK instructions cause the condition code in the current PSW to indicate the fact that the clock is invalid. Error logging is required as a result of clock failure.

Hard machine check interrupts on the System/370 Model 165 are as follows:

- System Damage - This interrupt occurs if PSW bit 13 is on. The SD bit in the stored machine check interrupt code (bit 0) is used to indicate that an error occurred during the execution of the instruction indicated by the machine check old PSW. The error was either a multiple-bit processor storage failure, a storage protection failure, a CPU error that was unretryable, or a CPU error that could not be corrected by the CPU retry hardware.

  If a multiple-bit processor storage failure caused the interrupt, the SE bit in the stored machine check code (bit 16) is on also, and the address of the failing storage area is indicated in the extended logout area. A storage protection failure is indicated by the KE bit (bit 18).

  A logout to the CPU extended logout area also occurs. For a CPU error, an extended logout occurs on the first and seventh retry. Error logging and the execution of recovery procedures are required after this type of interrupt.

Two modes of system operation will be used: full recording mode and quiet, or nonrecording, mode. In full recording mode, all machine check interrupt types cause an interrupt to be taken and logouts to occur. This will be the normal mode of Model 165 operation. In quiet mode, all or certain soft machine check interrupts are disabled. Quiet mode can be used to permit system operation without error recording for all or certain soft errors when a large number of transient (correctable) errors are occurring. It can also be used to allow Model 165 operation under the control of an operating system without Model 165 machine check handling routines included.

A hard stop status and a hard stop bit have been defined for the Model 165. The hard stop bit is located in control register 14 with the other mask bits discussed. If a hard stop condition occurs, the Model 165 system ceases all operations immediately without the occurrence of a logout to the fixed area. Hard stop is initiated by hardware rather than by programming.

When enabled, a hard stop occurs if a System Damage hard machine check type of error is detected when all machine check interrupts are disabled or during the processing of a previous machine check error.

Implementation of a hard stop prevents system operations from continuing when the nature of the machine malfunction prevents the system from presenting meaningful status data.

The state of the Model 165 after IPL or a system reset is:

1. CPU retry and an extended logout are disabled when a CPU error occurs. A hard machine check interrupt results on any CPU error, including those caused by a buffer malfunction.

2. Recovery reports are disabled. Therefore, single-bit processor storage error corrections do not cause a soft machine check interrupt.

3. External damage reports are enabled. A double- or multiple-bit processor storage failure associated with an I/O operation and damage to the time of day clock cause a machine check interrupt.

4. PSW bit 13 is normally enabled by the IPL PSW (it is disabled by system reset) so that System Damage machine checks (unretryable CPU failures, unsuccessfully retried CPU errors, storage protection failures, and multiple-bit processor storage errors associated with the CPU) cause a hard machine check interrupt.

5. Hard stop is enabled.

6. CPU extended logout is enabled and control register 15 points to location 512 as the beginning of the CPU extended logout area.

7. Channel logouts are disabled.

A machine check situation in a Model 65 or 75 results from hardware detection of a machine malfunction or of a parity error. Bad parity can occur in main storage, in local storage, in a register, in an adder, etc. Error correction hardware is not included in these two models. If the machine check mask in the current PSW (bit 13) is enabled, a machine check on Models 65 and 75 causes an interrupt and a diagnostic scan-out occurs, starting at location 128.

If recovery management support (RMS) is included in an OS control program for the Model 65, the MCH routine gains CPU control to record the error and to attempt corrective procedures after the machine check interrupt. Model 65 MCH performs programmed instruction retry for certain instructions only. Alternately, if an SER routine is present, the error is only logged, since a retry of the failing operation is not provided by this routine.

RECOVERY MANAGEMENT SUPPORT (RMS) - OS MFT AND MVT

MCH and CCH Routines

RMS for the Model 165 consists of extensions to the facilities offered by RMS routines currently provided for Models 65 and up. The two RMS routines, machine check handler (MCH) to handle machine check interrupts and channel check handler (CCH) to handle certain channel errors, will be included automatically in MFT and MVT control programs generated for the Model 165.

The two primary objectives of RMS are (1) to reduce the number of system terminations that result from machine malfunctions and (2) to minimize the impact of such incidents. These objectives are

accomplished by programmed recovery to allow system operations to continue whenever possible and by the recording of system status for both transient (corrected) and permanent (uncorrected) hardware errors.

MACHINE CHECK HANDLER: During IPL of a control program containing Model 165 RMS routines, machine check mask bits are enabled, and control register values are set to permit all machine check interrupts and logouts to occur.

MCH receives control after the occurrence of both soft and hard machine check interrupts. When a soft machine check occurs (successful CPU retry, single-bit processor storage error corrected, time of day clock damage, or multiple-bit processor storage error during an I/O operation), MCH formats a recovery report record to be written in the system error recording data set SYS1.LOGREC. This record contains pertinent information about the error, including pertinent data from the logout areas, an indication of the recovery that occurred, identification of the job, job step, and program involved in the error, the date, and the time of day. The operator is informed of successful CPU retries, single-bit processor storage corrections, and an error in the time of day clock.

MCH performs an additional function when a CPU retry was necessary because of a buffer malfunction. When an error occurs in the buffer, as indicated in the extended logout area, MCH updates a programmed buffer error counter. After a certain number of buffer errors occur, the entire high-speed buffer is disabled and MCH notifies the operator of this fact. The operator can allow the system to continue running in degraded mode, if necessary. All CPU fetches are then made directly to processor storage, bypassing the buffer. Alternately, the operator can terminate system operations and request CE diagnosis and repair of the buffer.

Prior to relinquishing CPU control, MCH determines whether or not an automatic mode switch from recording mode to quiet mode should take place if a CPU retry or an ECC correction recovery has just occurred. The determination of whether to switch to nonrecording (quiet) mode is made on the basis of the number of soft machine checks of a specific type that occur during system operation. Error count thresholds are maintained separately for successful CPU retry and successful processor storage single-bit error corrections. The IBM-supplied threshold values can be altered when the control program is generated.

MCH switches the system to quiet mode for either ECC corrections only (the DIAGNOSE instruction is used to change the ECC mode bit from full recording to quiet mode) or for both CPU retry and ECC corrections (the System Recovery mask bit is disabled). Mode switching occurs if the number of soft machine checks that occur during system operation exceeds the specified error count threshold for that type (or if SYS1.LOGREC is full). The operator is informed of the mode switch and can switch back to recording mode at any time thereafter.

Mode switching is implemented to attempt to prevent SYS1.LOGREC from being filled with recovery reports when a recurring correctable error condition exists that would cause many reports to be generated.

When a System Damage hard machine check occurs (uncorrectable or unretryable CPU error, multiple-bit processor storage error, or a storage protect key failure), MCH determines whether the error is one that is correctable by programming. A multiple-bit processor storage error or a storage protect key failure associated with CPU processing causes control to be given to the repair portion of the program damage assessment and repair (PDAR) routine of MCH. PDAR can repair damaged control program storage areas by loading a new copy of the affected module if the module is marked reentrant and refreshable (it has been

reorganization, and all processing programs that access the BDAM data set, is necessary if a relative track or actual address reference is used and fewer (or more) 3330 tracks are allocated to the data set than before.

Subject to the restrictions stated in Section 10:05 and those indicated for BDAM data sets, existing executable processing programs can be used without change to process data sets that have been transferred to 3336 packs. Nothing need be done to job control for these programs if the cataloged procedures supplied with the language translators are used, as long as the 3330 facility is specified as a SYSDA device at system generation. Otherwise, job control statements must be changed to request 3330 devices and, optionally, any data set characteristic changes, such as block size. RPS support, as described previously, is provided automatically.

User-written programs that use the EXCP level to access disk data sets that have been transferred to 3336 packs may have to be modified to reflect the characteristics of the data set on the 3336: a different number of records per track, a different number of tracks per cylinder, etc. All 2311 and 2314 CCW lists will operate on 3330 facilities except those that are device or channel time dependent and those that use the file scan commands, which are not available on the 3330. User-written 2311 or 2314 error routines will not execute correctly and must be modified. RPS commands have to be added by the user if this support is desired for programs that use EXCP. (Note that the XDAP macro will include support of RPS commands.)

Data sets currently located on 2301 Drum Storage can be placed on 2305 modules by means of a data set utility program. Unit specification in the job control statements of existing programs that will access the 2305 module instead of the 2301 drum must be changed. Also, to reflect the smaller capacity of a 2305 track, it may be necessary to alter the block size used. The latter can also be done via job control statement alterations (without program reassembly) unless block size was specified in the program itself or record length is larger than 2305 track size.

60:10 PLANNING OPTIMAL SYSTEM PERFORMANCE, USING BLOCK MULTIPLEXER CHANNELS AND RPS DEVICES

Block multiplexing, rotational position sensing, and multiple requesting provide the user with another tool that can improve total system throughput. However, the effectiveness of this tool for a given installation depends largely on proper planning for its use. This section indicates how to use block multiplexer channels and RPS devices effectively.

The guidelines outlined indicate how best to configure a system with rotational position sensing devices, how to plan job scheduling, and what to consider when determining disk data set characteristics. Explanations follow the statement of each guideline.

All guidelines presented are not necessarily practical for all users. Each item should be evaluated in terms of the processing requirements and hardware configuration of an installation.

SYSTEM CONFIGURATION AND GENERATION

System configuration and generation guidelines are as follows:

1. Multiple 3330 facilities should be placed on a single block multiplexer channel.

Performance improvement occurs (1) as a result of overlapping the rotational positioning time of disk devices and (2) because more I/O requests can be initiated in a given period of time, since the channel is free more often. When many disk devices are active concurrently on a block multiplexer channel, there is more potential for such overlap.

2. Direct access devices with RPS should be placed on separate channels from I/O devices without RPS. Alternate possibilities are as follows:

   a. If it is necessary to place non-RPS devices on the same block multiplexer channel with RPS devices, give first choice to non-RPS devices with a buffered control unit, such as the 2540 Card Read Punch and the 1403 Printer. These devices disconnect from a block multiplexer channel during the relatively long mechanical portion of their cycle, thereby freeing the channel for other operations.

   b. Tape units should not be placed on a block multiplexer channel with RPS devices unless absolutely necessary, because channel disconnection does not occur during any of their channel operations. If this is not possible, try to plan job scheduling to avoid having jobs using tape units and jobs using RPS support active on a block multiplexer channel at the same time. If this is not feasible, try to assign very low-activity data sets to these tape units.

      A device without channel disconnect capability can monopolize the block multiplexer channel for relatively long periods of time, thereby preventing (1) the initiation of other I/O operations on the channel and (2) the reconnection and completion of disk RPS channel programs already in operation on the channel. For example, a direct access device without RPS retains use of the channel during its search operations as well as during its reads and writes. If the device is a 2314 and block size is half a track, the channel is busy for 25 ms on the average (12.5 ms average rotational delay plus 12.5 ms read/write) for each I/O operation started for the non-RPS 2314 facility. Even if the block size used is relatively small, the channel can still be monopolized by the non-RPS device if there is high activity on the device.

3. The 2305 facility normally should not be placed on a block multiplexer channel with other devices.

   Exclusive use of a channel insures optimum performance of the 2305 facility as a system residence device.

The following should be noted in regard to specification of priority and ordered-seek I/O request queuing options for RPS devices at OS system generation. The priority queuing option insures priority I/O request initiation for the device, but because of first-come, first-served handling of I/O operations on the block multiplexer channel, this option does not insure that priority device channel programs will complete sooner than other RPS channel programs that were started later on the channel. However, the objective of specifying the ordered-seek queuing option (minimization of arm movement on a disk drive) can still be achieved when RPS is used.

SECTION 70: COMPARISON TABLE OF HARDWARE AND OS FEATURES FOR SYSTEM/360 MODEL 65 AND SYSTEM/370 MODEL 165

This table has been included for quick reference. It compares the hardware features and OS support of Models 65 and 165. A dash (-) in a programming system column indicates that the hardware feature does not require programmed support.

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
I. CPU				
A. Internal performance times Model 65	1	2 to 5	-	-
B. Instruction set				
1. Decimal and binary arithmetic	Standard	Standard	All languages except FORTRAN support of decimal arithmetic	Same
2. Floating-point arithmetic	Standard	Standard	All languages except RPG	Same
3. Extended precision floating-point	Not available	Standard	Mnemonics provided by Assemblers F and H. Subroutines provided in FORTRAN H	Subroutines and instruction support provided by same languages translators.
4. New general purpose instructions (MOVE LONG, COMPARE LONG, etc.)	Not available	Standard	-	Assembler Language mnemonics for user use
5. High-speed multiply	Not available	Optional	-	-
6. STORE CPU ID, STORE CHANNEL ID, HALT DEVICE instructions	Not available	Standard	-	Used by control program
C. Overlap of instruction fetching and preparation with execution	Instruction unit normally prepares one instruction at a time.	Instruction unit can process several instructions while execution unit executes one instruction.	-	-
	Imprecise interrupts occur only for storage protection violations.	Imprecise interrupts occur on an attempt to store at an invalid processor storage address or a storage protected location.		
D. Interval timer	Standard 16.6 ms resolution	Standard 3.33 ms resolution	Supported	Supported except for time of day requests
E. Time of day clock	Not available	Standard	-	Supported for time of day requests
F. CPU retry (hardware)	No	Standard	-	-

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
G. Machine check interrupt	One level of machine check provided for all machine errors. One machine check mask.	Occurs on both soft (corrected) and hard (uncorrected) errors. Four types are defined. Additional masking is provided.	SER0, SER1, and MCH routines. (See item VI for functions provided.)	MCH routine. (See item IV for functions provided.)
H. Size of fixed storage area in lower core	328 bytes including CPU and channel logouts	1512 bytes, reducible to 512 if the extended logout area is moved	Logout data processed by MCH or SER routine	Logout data processed by MCH
I. Control logic	Microprogram in ROS	Microprogram in both ROS and WCS	-	-
J. CPU cycle time	200 nanoseconds 8-byte data path	80 nanoseconds 8-byte data path	-	-
K. Direct Control	Optional	Standard	Not supported	Not supported
L. Compatibility features (all are optional and mutually exclusive)	1. 7070/7074 2. 7080 (for both 705 and 7080 emulation) 3. 709/7040/7044/7090/7094/7094II	1. 7074 (includes 7070 and 7074) 2. 7080 (for both 705 and 7080) 3. 709/7090/7094/7094II (does not include 704, 7040, or 7044)	Three stand-alone emulators are provided.	Three integrated emulator programs provided, one for each 7000 series compatibility feature
M. Multiprocessor systems	1. Multisystem optional feature permits interconnection of two Model 65s. Main storage is shared (512K or more). Direct control is required.	1. A multisystem feature is not available.	1. OS MP/65 supports Model 65 multiprocessor systems (MVT).	-
	2. The support or main processor in an ASP configuration can be a Model 65. Two or three systems are connected via a Channel to Channel Adapter.	2. A Model 165 can be a a support or main processor in an ASP configuration.	2. PCP, MFT, or MVT can be used for a Model 65.	2. MFT or MVT can be used for a Model 165.

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
II. STORAGE				
A. Processor storage sizes	256K 512K 768K 1024K	- 512K - 1024K 1536K 2048K 3072K	All are supported.	All are supported.
B. Processor storage cycle	750 nanoseconds (for 8 bytes). Two-way interleaving of sequential accesses other than by the channels is provided.	2 microseconds. Storage is 4-way, doubleword interleaved for CPU and channel requests. 32 bytes can be obtained every 2 microseconds.	-	-
C. Processor storage validity checking	Parity checking by byte. No hardware error correction is provided.	ECC checking on a doubleword. Single-bit errors are corrected by hardware.	-	-
D. Processor storage reconfiguration	Not provided for uniprocessor systems	Malfunctioning storage boxes can be configured out of the operational system.		
E. Byte boundary alignment permitted for operands of nonprivileged instructions	No	Standard	-	Programmers can byte-align in Assembler, PL/I, ANS COBOL, and FORTRAN programs.
F. Storage and fetch protection	Standard	Standard	Storage protection is supported.	Same as Model 65
G. 2361 Core Storage	Optional Up to 8 million bytes can be attached.	Cannot be attached	Hierarchy support is provided.	-
H. High-speed buffer storage	Not provided	8K is standard. 8K more can be added	-	-
1. Cycle time		80 nanoseconds	-	-
2. CPU-to-buffer fetch time		160 nanoseconds for 8 bytes; a request can be initiated every 80 ns cycle.	-	-

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
III. CHANNELS				
A. Total number per system	Up to 7	1. Up to 7 2. Up to 12 with Extended Channels special feature	Supported	Up to 12 channels supported
B. 2870 Multiplexer Channel	One or two can be attached.	Same as Model 65	Supported	Supported
C. 2860 Selector Channel (1.3 MB)	A maximum of 6 can be attached.	A maximum of 6 can be attached.	Supported	Supported
D. 2880 Block Multiplexer Channel (1.5 MB). Two-Byte Interface feature permits a 3.0 MB data rate.	Cannot be attached	A maximum of 6 can be attached without the Extended Channels feature, a maximum of 11 with this feature.	-	Supported in selector or block multiplexer mode
E. Aggregate channel data rate	In excess of 4 MB with one 2870 and six 2860s	In excess of 9 MB with one 2870 and eleven 2880 channels		
F. Channel retry data provided	Yes in I/O logout	Yes in I/O logout	Channel retry provided by CCH routine	Channel retry provided by CCH routine
G. Channel-to-Channel Adapter on 2860	Optional	Optional	-	-
IV. OPERATOR CONSOLE DEVICES	1. 1052 Printer-Keyboard (optional) 2. Second 1052 Printer-Keyboard is optional. 3. A 2250 Display Unit and a remote 2150 Console are optional. 4. Other devices can be used as primary and secondary consoles a. 1052 Printer-Keyboard b. 2260 Display Station c. 2250 Display Unit	1. Stand-alone 3066 System Console is required. It includes: a. A CRT-Keyboard combination for operator/system communication b. An indicator viewer c. A microfiche document viewer d. A processor storage configuration plugboard e. A system activity meter	1. Supported as the primary console 2. MCS option supports multiple consoles. 3. Supported as the primary console 4. Additional consoles supported by MCS and DIDOCS options	1. Supported as the primary system console a. Supported by the MCS options and DIDOCS

Features	System/360 Model 65	System/370 Model 165	OS - Model 65	OS - Model 165
		f. A device for loading WCS and microdiagnostics 2. Optionally, other devices can be used as secondary consoles as listed for the Model 65.		2. Supported by MCS and DIDOCS options
V. I/O DEVICES				
A. 3211 Printer	Yes	Yes	Supported	Supported
B. 3330 facility	No	Yes on a 2880	-	Supported including RPS and multiple requesting
C. 2305 facility Models 1 and 2	No	Yes on a 2880	-	Supported including RPS and multiple requesting
VI. PROGRAMMED RECOVERY ROUTINES FOR HARDWARE ERRORS				
A. SER0 in control program	Optional (Either SER0, SER1, or MCH must be included.)	-	Logs status data and terminates system operation	Not provided
B. SER1 in control program	Optional (Either SER0, SER1 or MCH must be included.)	-	Logs status data and attempts to terminate affected task and continue system operation	Not provided
C. Recovery management support (RMS) - MCH and CCH in program control	Optional	Required	1. Programmed instruction retry is attempted for some instructions. 2. Provides repair of refreshable control program and processing program modules damaged by a hardware error.	1. Status data logged after both successful and unsuccessful hardware retry of CPU errors 2. Single-bit errors are corrected by ECC hardware. Status data is logged by MCH. Repair of refreshable control program modules damaged by double- and multiple-bit processor storage errors and CPU errors is provided.

INDEX

READER'S COMMENT FORM

A Guide to the IBM System/370
Model 165

GC20-1730-0

Please comment on the usefulness and readability of this publication, suggest additions and deletions, and list specific errors and omissions (give page numbers). All comments and suggestions become the property of IBM. If you wish a reply, be sure to include your name and address.

COMMENTS

fold

fold

fold

fold

- Thank you for your cooperation. No postage necessary if mailed in the U.S.A. FOLD ON TWO LINES, STAPLE AND MAIL.

YOUR COMMENTS PLEASE...

Your comments on the other side of this form will help us improve future editions of this publication. Each reply will be carefully reviewed by the persons responsible for writing and publishing this material.

Please note that requests for copies of publications and for assistance in utilizing your IBM system should be directed to your IBM representative or the IBM branch office serving your locality.

fold fold

FIRST CLASS
PERMIT NO. 1359
WHITE PLAINS, N. Y.

BUSINESS REPLY MAIL
NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

POSTAGE WILL BE PAID BY...

IBM Corporation
112 East Post Road
White Plains, N. Y. 10601

Attention: Technical Publications

fold fold

A Guide to the IBM System/370 Model 165 Printed in U.S.A. GC20-1730-0

IBM

International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N.Y. 10601
[USA Only]

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
[International]

IBM **Technical Newsletter**

Re: Order No. GC20-1730-0

This Newsletter No. GN20-2227

Date July 14, 1970

Previous Newsletter Nos. None

A GUIDE TO THE IBM SYSTEM/370 MODEL 165

This Technical Newsletter provides replacement pages for the subject manual. Pages to be inserted and/or removed are listed below.

5,6	65,66
11,12	69,70
21–24	77–80
35–38	97–100
53–56	105–110

A change to the text or a small change to an illustration is indicated by a vertical rule to the left of the change. A changed or added illustration is denoted by the symbol • to the left of the caption.

Summary of Amendments

Minor additions and corrections have been made throughout the text.

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

IBM Corporation, Technical Publications Dept., 112 E. Post Road, White Plains, N. Y. 10601

Printed in U.S.A.

OPTIONAL FEATURES

The following optional features can be field installed unless otherwise noted:

- High-Speed Multiply (plant installation)
- 709/7090/7094/7094II Compatibility Feature (plant installation)
- 7070/7074 Compatibility Feature (plant installation)
- 7080 Compatibility Feature (plant installation)
- Buffer Expansion feature for inclusion of a 16K buffer
- Channel-to-Channel Adapter - on 2860 only
- Extended Channels (up to 12 addressable channels)
- Remote Operator Console Panel Attachment

Note: Compatibility features are mutually exclusive.

SECTION 20: I/O DEVICES

20:05 I/O DEVICE SUPPORT

Most presently announced I/O devices that can be attached to System/360 Models 65 and above can be attached to the System/370 Model 165. The following I/O devices are not included in a standard Model 165 configuration:

1017	Paper Tape Reader
1018	Paper Tape Punch
1052	Printer-Keyboard, Models 3, 5, 6, and 8
1231	Optical Mark Page Reader
1259	Magnetic Character Reader
1285	Optical Reader
1404	Printer
1412	Magnetic Character Reader
1418	Optical Character Reader
1428	Alphameric Optical Reader
1445	Printer
1827	Data Control Unit (for attachment of 1800 system analog and/or digital control units to the Model 165)
2302	Disk Storage
2415	Magnetic Tape Unit and Control
2671	Paper Tape Reader
7340	Hypertape Drive
7772	Audio Response Unit

New I/O devices for the Model 165 are:

- The 3330 facility - attaches only to a block multiplexer channel
- The 2305 facility Models 1 and 2 - attaches only to a block multiplexer channel
- The 3211 Printer - attaches to any Model 165 channel

The 3330 and the 2305 facilities represent significant advancements in direct access device technology. They provide greater online data capacity and faster data rates and access than heretofore, as well as expanded error correction features. Both have rotational position sensing and multiple requesting as standard features.

The 3330 represents the latest generation of direct access device with removable, interchangeable disk packs. The 2305 facility is a major extension of the nonremovable, high-speed, fixed-head direct access storage concept.

The major new characteristics of the 3330 and 2305 facilities and the 3211 Printer are discussed in the following subsections.

20:10 3330 DISK STORAGE AND 3830 STORAGE CONTROL

The 3330 facility is a modular, large-capacity, high-performance direct access storage subsystem. The 3330 facility consists of 3830 Storage Control and from one to four 3330 Disk Storage modules. A 3330 module contains a pair of independent disk storage drives as shown in Figure 20.10.1. The new, removable 3336 Disk Pack is used for data storage. Usage meters are contained in the 3830 control unit and in each 3330 module.

- End-of-volume (EOV) routines will support concatenation of data sets residing on RPS and on non-RPS devices. The control program will ensure that an RPS access method is used for drives with the feature and that a non-RPS method is used for drives without it.

- Any system utility, data set utility, or IBM-supplied processing program (such as a language translator) that uses the sequential access methods will support RPS. In addition, IEHMOVE, IEBCOPY, and the initialize/dump/restore utility (IEHDASDR) will include RPS support for 3330 facilities.

- The Sort/Merge program product supports RPS for 3330 intermediate work devices.

- Where appropriate, RPS commands for access to SYSRES data sets will be supported by:

  Data set catalog routines
  Direct access device space management (DADSM) routines (for DSCB processing)
  STOW, BLDL, and FIND processing of program library (PDS) directory entries
  OPEN/CLOSE/EOV processing of JFCB's in the job queue
  Routines that access the job queue

  Note that RPS command support is not provided for:

  Access to TSO swap data sets
  Program fetch
  The telecommunications access method (TCAM) for message queues
  The stand-alone disk initialization and alternate track assignment routines (DASDI and IEHATLAS)

The 2305 facility. The 2305 facility will be supported as an I/O device for most of the same functions as is 2301 Drum Storage. Note that specification of exchange buffering for a QSAM data set on either a 2305 or 2330 facility results in a default to simple buffering.

Rotational position sensing for 2305 facilities will be supported as discussed for the 3330 facility except that an arm positioning seek is not required on 2305 modules and the 2305 is not supported as an intermediate work unit by Sort/Merge programs. In addition, the 2305 is not supported by ASP, HASP, or for TCAM message queues. Multiple requesting will be handled by the I/O supervisor, which will initiate up to seven channel programs on one 2305 module at a time. (The eighth logical device address is reserved for control purposes.) Multiple I/O operations to the same data set will be initiated if the data set is not being sequentially processed. (One operation must complete successfully before the next can be initiated during sequential processing.)

3211 Printer with tapeless carriage. The 3211 Printer, with or without the 18 additional print positions, will be supported by QSAM and BSAM for exactly the same functions as is the 1403 Printer and by ASP and HASP. The control program will handle loading of the forms control buffer (FCB) with carriage control images. This support is similar to that provided for Universal Character Buffer (UCB) loading.

The user can define one or more default FCB images at system generation time. Two IBM-supplied default images are included automatically. All other FCB images to be used must be defined by the user and placed in SYS1.SVCLIB, as is the case with UCB images. User-supplied default images must be identified as defaults. The FCB image to be used by a processing program can be specified in the 3211

Printer DD statement included for the job step and will be loaded into the FCB by the control program.

The FCB image currently loaded can also be changed by the programmer during execution of the processing program by use of an Assembler Language macro. If the DD statement does not specify an FCB image and the image currently loaded is not one of the defaults specified at system generation, the operator is requested to specify the FCB image to be used.

The FCB image (and the UCS character image) to be loaded for a 3211 Printer used by an output writer can be indicated in the output writer procedure or in the START output writer command issued by the operator. FCB and UCB images can also be specified in the SYSOUT DD statement for a data set that is to be printed by the output writer, and they will be loaded into the 3811 control unit prior to the printing of the data set.

Any time the FCB parameter is used, as described above, the user can specify that operator verification of forms alignment is to be requested by the control program via a console message when the buffer is loaded. The operator must respond to this message.

The 3211 error recovery routine will retry a print operation after a parity check occurs in the UCB or print line buffer if QSAM is used and if three I/O buffers are provided for the printer data set. When the operation is retried, the 3811 control unit insures that only the print positions that did not print correctly the first time are reprinted.

Operator console. The alphameric keyboard and CRT will be supported as the primary operating system console. The multiple console support (MCS) and device-independent display operator console (DIDOCS) options of OS must be included in the operating system to support this console.

IBM **Technical Newsletter**

Re: Order No. GC20-1730-0

This Newsletter No. GN20-2277

Date November 20, 1970

Previous Newsletter Nos. GN20-2227

A GUIDE TO THE IBM SYSTEM/370 MODEL 165

This Technical Newsletter provides replacement pages for the subject manual. Pages to be inserted and/or removed are listed below.

Preface and Contents	51.12, 52
3, 4	55, 56
25 – 28	89 – 90.1
35 – 36.1	101, 102
51, 51.1	109 – 111
51.2 – 51.11	

A change to the text or a small change to an illustration is indicated by a vertical rule to the left of the change. A changed or added illustration is denoted by the symbol • to the left of the caption.

Summary of Amendments

Additions and changes have been made to include information about the 3803/3420 Magnetic Tape Subsystem.

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

IBM Corporation, Technical Publications Dept., 112 E. Post Road, White Plains, N. Y. 10601

Printed in U.S.A.

A Guide to the IBM System/370 Model 165 Printed in U.S.A. GC20-1730-0

International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, New York 10601
(USA only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)