

Systems

A Guide to the IBM 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex of System/370

This guide presents hardware, programming systems, and other pertinent information describing the significant new features and advantages of the IBM 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex. Knowledge of a System/370 processor that has EC mode and dynamic address translation capabilities is assumed. This guide is intended to acquaint the reader with the components of 3033 configurations and to be of benefit in installation planning.

Associated with this guide are three optional supplements describing programming systems for the 3033 Processor Complex that support a virtual storage environment. Each supplement has its own form number and must be ordered individually, if desired. The optional supplements are:

- *OS/Virtual Storage 1 Features Supplement (GC20-1752)*
- *OS/Virtual Storage 2 Single Virtual Storage (SVS) Features Supplement (GC20-1753)*
- *Virtual Machine Facility/370 Features Supplement (GC20-1757)*



Fifth Edition (April 1979)

This edition is a major revision obsoleting GC20-1859-3. A discussion of the 3033 Attached Processor Complex configuration has been added and the new processor storage sizes have been included. Various other changes have been made as well.

This guide is intended for planning purposes only. It will be updated from time to time; however, the reader should remember that the authoritative sources of system information are the system library publications for the 3033, its associated components and its programming support. These publications will first reflect any changes.

Requests for copies of IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

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, Dept. 824, 1133 Westchester Avenue, White Plains, New York 10604. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

© Copyright International Business Machines Corporation 1977, 1978, 1979

PREFACE

It is assumed that the reader of this publication is familiar with System/370 architecture, features, and I/O devices that are common among System/370 processors in which EC mode is implemented. Knowledge of virtual storage and virtual machine concepts and advantages is also assumed. Such information can be found in other guide series publications, such as A Guide to the IBM System/370 Model 168 for System/360 Users (GC20-1787), and system library publications.

This guide identifies the major facilities of the 3033 Processor Complex in Sections 05 through 20 but detailed explanations and advantages are given only for features or implementations that are different from those of most other System/370 processors with dynamic address translation capability. Additional, more detailed information regarding 3033 Processor Complex hardware and programming systems support can be found in system library publications.

The 3033 Multiprocessor Complex and 3033 Attached Processor Complex, which are tightly coupled multiprocessing configurations, are also discussed. This guide does not assume knowledge of multiprocessing concepts, hardware, or advantages, all of which are covered.

The total guide for the 3033 Processor Complex consists of this base publication (Sections 01 through 70), which covers 3033 Processor, 3036 Console, and 3037 Power and Coolant Distribution Unit hardware, and from one to three optional supplements (Sections 90, 100, and 110). The optional supplements describe the facilities of three of the IBM programming systems that support a virtual storage environment using the dynamic address translation hardware of the 3033 Processor. Each optional supplement has its own unique form number and each supplement desired must be ordered separately and inserted in this base publication, which is distributed without the automatic inclusion of any optional supplements.

The following optional supplements can be inserted in this base publication:

- OS/Virtual Storage 1 Features Supplement (GC20-1752) - assumes knowledge of OS MFT
- OS/Virtual Storage 2 Single Virtual Storage (SVS) Features Supplement (GC20-1753) - assumes knowledge of OS MVT
- Virtual Machine Facility/370 Features Supplement (GC20-1757) - does not assume knowledge of CP-67/CMS

All optional supplements also assume knowledge of virtual storage, dynamic address translation, and other new 3033 Processor and 3036 Console features as described in this base publication or appropriate system library documents. However, no optional supplement requires knowledge of the contents of any other optional supplement.

This base publication, as well as each optional supplement, begins with page 1 and includes its own table of contents and index. The base publication or supplement title is printed at the bottom of each page as a means of identification.

The optional programming systems supplements contain System/370 processor-independent information, unless otherwise noted, and are designed to be included in the guides for the 3031, 3032, and 3033 Processors and Models 135, 138, 145, 148, 158, and 168 of System/370, as shown in the table that follows.

Base Publications	Supplements			
	DOS/VS Features Supplement (GC20-1756)	OS/VS1 Features Supplement (GC20-1752)	OS/VS2 SVS Features Supplement (GC20-1753)	VM/370 Features Supplement (GC20-1757)
A Guide to the IBM System/370 Model 135 (GC20-1738)	X	X		X
A Guide to the IBM System/370 Model 138 (GC20-1785)	X	X		X
A Guide to the IBM System/370 Model 145 (GC20-1734)	X	X	X	X
A Guide to the IBM System/370 Model 148 (GC20-1784)	X	X		X
A Guide to the IBM System/370 Model 158 for System/370 Model 155 Users (GC20-1754)	X	X	X	X
A Guide to the IBM System/370 Model 158 for System/360 Users (GC20-1781)	X	X	X	X
A Guide to the IBM System/370 Model 168 for System/370 Model 165 Users (GC20-1755)		X	X	X
A Guide to the IBM System/370 Model 168 for System/360 Users (GC20-1787)		X	X	X
A Guide to the IBM 3031 Processor Complex and Attached Processor Complex of System/370 (GC20-1854)	X	X	X	X
A Guide to the IBM 3032 Processor Complex of System/370 (GC20-1858)		X	X	X
A Guide to the IBM 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex of System/370 (GC20-1859)		X	X	X

A Guide to the IBM 3033 Processor Complex, AP Complex, and MP Complex

CONTENTS

Base Publication Sections (Sections 01 through 70)

Section 01:	Highlights of the 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex	1
Section 05:	3033 Processor Complex Physical Components and Technology	6
	Physical Components	6
	Motor Generator Set	7
	Cooling	7
	Technology	8
Section 10:	3033 Processor Architecture Design and Functional Components	9
10:05	Architecture Design and Processor Elements	9
	Architecture Design	9
	Processor Elements and Features	10
10:10	Instruction Preprocessing and Execution Functions	14
	Instruction Preprocessing Function	14
	Execution Function	16
10:15	Processor Storage and the Processor Storage Control Function	24
	Processor (Main) Storage	24
	Processor Storage Control Function	27
10:20	Channels	36
	Channels and Directors	36
	Channel Bus Controller	44
	I/O Error Handling	44
10:25	Maintenance and Retry Function	46
	Trace Unit	46
Section 20:	The 3036 Console for the 3033 Processor Complex	49
20:05	Components, Functions, and General Operation	49
	Components	49
	Functions and General Operation	51
20:10	Control Panel	52
20:15	CRT Display and Keyboard	56
20:20	Diskette Drive and Diskettes	57
20:25	Display Frames	59
	Operator Frames	60
	Service Frames	66
	Indicator/Logout Frames	68
20:30	Console Configuration	68
	Types of Configurations	68
	Normal Console Configuration	70
	Maintenance Mode Console Configurations	72
	Console Recovery Facilities	75
20:35	Serviceability Features	76
	Voltage Monitoring	76
	Fault Locating Tests	76
Section 30:	The 3033 Multiprocessor Complex	78
30:05	General Description	78
	Introduction	78
	General Definition of Multiprocessing	78
	The 3033 Multiprocessor Complex Configuration	80
	Advantages of Tightly Coupled Multiprocessing Configurations	91

	Advantages of Loosely Coupled Multiprocessing Configurations	96
30:10	3033 Multiprocessing Architecture and Hardware	98
	3033 Processor Complex and Multiprocessor Complex Hardware Differences	98
	Prefixing	99
	Processor Addressing	100
	Time-of-Day Clock	100
	Interprocessor Programmed Communication	101
	Interprocessor Hardware Communication	103
30:15	Recording and Diagnostic Programs	104
30:20	Planning Considerations	105
	Planning for Maximum System Availability	105
Section 40:	The 3033 Attached Processor Complex	108
40:05	General Description	108
	Components	108
	3033 Model A Processor	110
	3042 Attached Processor	113
	3038 Multiprocessor Communication Unit	113
	3036 Consoles	113
40:10	3033 Processor Complex and Attached Processor Complex Hardware Differences	118
40:15	Advantages of a 3033 Attached Processor Configuration	121
	Less Complex Operational Requirements	121
	Improved Resource Utilization	121
Section 50:	Programming Systems Support	123
50:05	General Description	123
50:10	MVS/System Extensions Program Product	123
50:15	VM/System Extensions and VM/Basic System Extensions Program Products	124
Section 60:	Highlights of the 3033 Processor Complex for Model 168 Users	126
	Price Performance Improvements	126
	Channel Enhancements	128
	Availability and Serviceability Features	128
	New Technology	129
Section 70:	Comparison Tables	130
70:05	Comparison Table of Hardware Features of System/370 Models 165, 165 II, and 168 (Models 1 and 3), and the 3033 Processor Complex	131
70:10	OS/VS Support of the 3033 Processor Complex	140
Index (Sections 01 through 70)		145
<u>Optional Sections</u> (See each supplement for detailed contents and index)		
Section 90:	OS/Virtual Storage 1 Features	151
Section 100:	OS/Virtual Storage 2 Single Virtual Storage (SVS) Features	153
Section 110:	Virtual Machine Facility/370 Features	155
<u>Note:</u>	This guide does not have a Section 80. DOS/Virtual Storage Features are discussed in the Section 80 supplement and 3033 Processors are not supported by DOS/VS.	

FIGURES (Sections 01 through 70)

05.1	Major components of the 3033 Processor Complex	6
05.2	Conceptual flow of the water cooling system in the 3033 Processor	8
10.05.1	Major elements in the 3033 Processor	11
10.10.1	Processor-dependent fixed storage locations	19
10.10.2	3033 Processor machine check code	20
10.15.1	Processor storage control function components and interface with other processor elements	28
10.15.2	64K buffer organization	30
10.15.3	Processor storage address format for buffer reference	31
10.15.4	TLB format	33
20.05.1	The 3036 Console	50
20.10.1	Control panel on the 3036 Console	53
20.25.1	The operator frame	61
20.25.2	The program frame	62
20.25.3	The alter/display frame	63
20.25.4	The configuration frame	63
20.25.5	The system activity frame	65
20.25.6	The TP link frame	67
20.30.1	The normal 3036 Console configuration	70
20.30.2	A maintenance mode 3036 Console configuration	73
20.30.3	Director servicing using a maintenance mode console configuration	74
30.05.1	The 3033 Multiprocessor Complex	81
30.05.2	The normal console configuration for a 3033 Multiprocessor Complex	86
40.05.1	Components of the 3033 Attached Processor Complex	109
40.05.2	The normal console configuration for the 3033 Attached Processor Complex	115

TABLES (Sections 01 through 70)

10.10.1	3033 Processor machine check interruptions	22
10.15.1	3033 Processor cycle and access times	26
10.20.1	Channel attachment and UCW mode for the more frequently used I/O devices that attach to the 3033 Processor	40
40.10.1	Comparison of hardware features in the Model U 3033 Processor, Model A 3033 Processor, and 3042 Attached Processor	120

SECTION 01: HIGHLIGHTS OF THE 3033 PROCESSOR COMPLEX, ATTACHED PROCESSOR COMPLEX, AND MULTIPROCESSOR COMPLEX

The 3033 Processor Complex consists of the 3033 Processor, 3036 Console, and 3037 Power and Coolant Distribution Unit. The 3033 Processor is a high-speed, large-scale, advanced function processor of System/370. It has a significantly higher internal performance than System/370 Models 165 and 168. It offers these users new levels of price performance, new function, and improvements in availability and serviceability without the necessity of reprogramming. All but a few of the features for the 3033 Processor are standard.

The 3033 Processor is a general purpose processor and offers high performance for both commercial and scientific applications. The 3033 Processor has hardware features and programming systems support, such as that for virtual storage and virtual machines, that are designed to facilitate application development and maintenance. In addition, a 3033 Processor Complex, its I/O devices, and its programming support can ease the expansion of data base and online data processing operations.

The 3033 Multiprocessor Complex is a tightly coupled multiprocessing configuration that consists of two 3033 Processors with multiprocessing hardware interconnected via the 3038 Multiprocessor Communication Unit, two 3036 Consoles, and two 3037 Power and Coolant Distribution Units.

The 3033 Attached Processor Complex is a tightly coupled multiprocessing configuration that consists of a 3033 Processor with multiprocessing hardware interconnected to a 3042 Attached Processor via the 3038 Multiprocessor Communication Unit, two 3036 Consoles, and two 3037 Power and Coolant Distribution Units. Like a 3033 Multiprocessor Complex, the attached processor configuration can execute two instruction streams (tasks) simultaneously, one in each processor.

Highlights of the 3033 Processor Complex are as follows.

- The 3033 Processor is capable of an instruction rate generally in the range of 1.6 to 1.8 times the System/370 Model 168 Model 3 with identical programs and configurations running under OS/VS2 Release 3. The upper end of the range is identified by jobs generally of the commercial and interactive mix. The lower end of the range represents an environment generally associated with scientific jobs, such as FORTRAN Compile and Go.
- The cycle time of the 3033 Processor is 57 nanoseconds.
- Basic control (BC) mode, for compatibility with System/360, and the System/370 extended control (EC) mode of processor operation are standard.
- The standard instruction set provides binary, decimal, and floating-point (including extended precision) arithmetic operations. It includes all the new System/370 instructions, that is, those provided for System/370 but not System/360 (see Section 10:05), as well as additional instructions that are part of the standard System/370 Extended Facility.
- The System/370 Extended Facility is standard. It provides new function. When supported by the MVS/System Extensions program product, this facility provides the means to (1) help reduce the time needed to execute certain frequently used supervisor functions of MVS, (2) increase the efficiency of dynamic address translation, and (3) improve processor availability through additional protection of certain processor storage locations that are vital to operating system availability.

- Separate instruction preprocessing and execution functions are implemented that provide overlap of instruction fetching, instruction decoding, operand fetching, and instruction execution to increase internal performance. A portion of the increased internal performance of the 3033 Processor versus the 3168 Processing Unit is the result of several improvements in instruction preprocessing and execution function design.
- Dynamic address translation (DAT) is a standard facility that can be made operative only when the 3033 Processor is operating in EC mode. It provides hardware translation of addresses during program execution. One virtual storage of up to 16 million bytes or multiple virtual storages of up to 16 million bytes each can be supported using DAT hardware. (The amount of virtual storage that can be efficiently supported by a 3033 Processor depends on the hardware configuration and job stream characteristics.)
- Channel indirect data addressing is a standard feature. Channel indirect data addressing enables the channels to access an I/O buffer that is contained in noncontiguous processor storage areas and is used when DAT mode is operative.
- An interval timer of 3.33 ms resolution and a time-of-day clock with a one-microsecond resolution are standard.
- A CPU timer and clock comparator are standard. The CPU timer provides an interval timing capability similar to that of the interval timer but it is updated every microsecond, as is the time-of-day clock. The clock comparator can be used to cause an interruption when the time-of-day clock passes a specified value. These items provide higher resolution timing facilities than the interval timer and enable more efficient timing services routines to be written.
- Program event recording (PER) is standard and can be made operative when the 3033 Processor is operating in EC mode. It is designed to be used as a problem determination aid. This feature includes hardware that can monitor the following during program execution: successful branches, the alteration of general registers, and instruction fetches from, and alterations of, specified areas of processor storage.
- A monitoring feature is standard and can be used to trace user-defined program events for the purpose of debugging or statistics gathering.
- The standard byte-oriented operands facility permits byte boundary alignment of the operands of nonprivileged instructions, making it unnecessary to add padding bytes within records or to blocked records for the purpose of aligning fixed- or floating-point data. Performance degradation for instructions that use unaligned data is greatly reduced in the 3033 Processor, when compared with the degradation that occurs in a Model 168, as a result of instruction preprocessing and execution function design changes. Only minimal degradation occurs in the 3033 Processor.
- Reloadable control storage is included in the 3033 Processor to contain the microcode for the instruction set and, during servicing operations, to contain fault locating microdiagnostics. Separate reloadable control storage is provided for the microcode for the channels.
- Instruction retry of most failing processor hardware operations is handled automatically by the hardware, without programming assistance.

- Storage features of the 3033 Processor are as follows:

A two-level storage system, consisting of fast, large-size processor (main) storage used as backing storage for a smaller, high-speed buffer storage, is implemented.

64K bytes of monolithic buffer storage with a 57-nanosecond cycle time are provided. The instruction processor function can fetch eight bytes from the buffer in 114 nanoseconds.

Processor storage sizes of 4M-, 8M-, 12M-, and 16M-bytes (M=1,048,576) are available. Processor storage is eight-way, doubleword interleaved and has a read/write cycle time of 285 nanoseconds for a doubleword. It is implemented using monolithic technology. Store and fetch protection are standard.

Error checking and correction (ECC) hardware that automatically corrects all single-bit processor storage errors and detects all double-bit and many multiple-bit errors is standard. More multiple-bit errors are detected in the 3033 Processor than in the Model 168 as a result of the implementation of a new ECC code.

- Channel features and I/O devices for the 3033 Processor are as follows:

Channels are packaged in groups and each group is controlled by a microprogrammed director. Channels are contained within the 3033 Processor but each channel group and its director has its own control storage, microprogram, and arithmetic logic unit so that channel operations interfere with instruction processing operations only when a channel and the instruction processor need access the same logical storage element simultaneously. A channel group can be serviced concurrently with normal processing and can be powered on and off separately from other channel groups and processor components.

Two channel groups, each of which consists of one byte multiplexer and five block multiplexer channels, are standard. Optionally, one additional channel group consisting of four block multiplexer channels or one byte and three block multiplexer channels can be installed for a total of 16 channels in the processor. A block multiplexer channel can operate at a data rate of up to 1.5 megabytes per sec (MB/sec). The five block multiplexer channels in a group are capable of achieving an aggregate data rate of 6.7 MB/sec when certain configuring rules are observed.

The optional Two-Byte Interface feature can be installed on the first block multiplexer channel in each channel group to increase its maximum data rate to up to 3 MB/sec. Optionally, one or two Channel-to-Channel Adapters can be installed in a 3033 Processor, one in each standard channel group.

Channel retry data is provided after channel errors so that error recovery routines can retry I/O operations. Significantly more logout data is provided after certain channel errors than for most other System/370 processors.

Most I/O devices that attach to the channels for a 3168 Processing Unit can be attached to the channels in a 3033 Processor.

- Space requirements for a 3033 configuration are significantly less than for a Model 168 configuration, since channels are contained within the 3033 Processor. In addition, the 3033 Processor requires less power, air cooling, and water cooling than a 3168 Processing Unit with its standalone channels because of the new channel design and new technology for logic and processor storage.

- The console for the 3033 Processor is the standalone 3036 Console, which provides significant new operational and servicing capabilities.

The 3036 Console provides the means for manually controlling 3033 Processor operations, controlling power on and off operations for processor complex components and subcomponents, loading microcode and microdiagnostics, configuring certain 3033 and 3036 components, and performing many diagnostic procedures.

The 3036 Console contains two physically and functionally separate, individually addressable operating stations. Each station contains a CRT display and keyboard, diskette drive for loading microcode and microdiagnostics, and microprogram-controlled console processor that controls the operation of the station. Either station can be used to control normal processing operations while the other is used for service support operations. Alternatively, both stations can be used for normal processing operations or service operations.

The dual station design permits one of several types of maintenance operations, including one remote diagnostic procedure, to be performed concurrently with normal processing operations. This dual station design also provides console backup.

Highlights of the 3033 Multiprocessor Complex are as follows:

- The IBM 3033 Multiprocessor Complex is capable of an instruction execution rate generally in the range of 1.6 to 1.8 times that of the 3033 uniprocessor with similar configurations and identical programs running under OS/VS2 MVS.
- All standard and optional features that are available for the 3033 Processor Complex are available for the 3033 Multiprocessor Complex.
- The features of the 3033 Multiprocessor Complex provide the advantages normally provided by a tightly coupled multiprocessing as compared with two uncoupled uniprocessor configurations: improved availability, less complex operational requirements, improved resource utilization, operational flexibility, improved growth options, and improved throughput possibilities.
- The design of the multiprocessing capability for a 3033 Multiprocessor Complex provides advantages over a Model 168 Multiprocessing System, such as enhanced channel recovery hardware for use after an unrecoverable processor failure, a more convenient configuration capability via elimination of a configuration control panel, and serviceability improvements (as for a uniprocessor configuration) provided by the 3036 Console.

Highlights of the 3033 Attached Processor Complex are as follows:

- The 3033 attached processor configuration contains two tightly coupled processors that are equivalent in function to the 3033 Processor for uniprocessor 3033 configurations except for the absence of any processor storage and channels in the 3042 Attached Processor. The 3033 and 3042, operating under the control of a single multiprocessing operating system, share all processor storage in the 3033 Processor and all I/O operations are performed by the 3033.
- The 3033 Attached Processor Complex is capable of an instruction execution rate generally in the range of 1.6 to 1.8 times that of a 3033 Processor Complex with similar configurations and identical programs running under OS/VS2 MVS.

- The 3033 Attached Processor Complex offers increased internal performance over a 3033 Processor Complex and provides advantages over two uniprocessor 3033 configurations, such as a single system image and improved resource utilization.

Basic support of the 3033 Processor Complex is provided by the following system control programs: OS/VS1, OS/VS2 MVS and SVS, and VM/370. The increased internal performance achieved when using these programming systems on a 3033 Processor versus a Model 168 can be of particular aid in expanding data base and online applications.

Support of the new function of the 3033 Processor, the System/370 Extended Facility, is provided by the MVS/System Extensions program product, which also offers several other performance enhancements for MVS users. The faster internal performance of the 3033 Processor and its large processor storage sizes combined with the support provided by MVS/System Extensions enables a 3033 installation to better utilize the benefits of MVS.

The VM/System Extensions program product is designed to improve the performance of VM/370 through the utilization of a resource manager, increase the performance of virtual machines in which a virtual storage programming system is executing, support an MVS operating system with the MVS/System Extensions program product installed executing in a virtual machine, and provide other new functions. The VM/Basic System Extensions program product provides a subset of the new functions offered by VM/System Extensions.

Tightly coupled multiprocessing support for the 3033 Multiprocessor Complex is provided by OS/VS2 MVS. The System/370 Extended Facility can be utilized in both processors during multiprocessor mode operations when the MVS/System Extensions program product is installed.

Tightly coupled multiprocessing support for the 3033 Attached Processor Complex is provided by OS/VS2 MVS and VM/370. The MVS/System Extensions program product supports the use of the System/370 Extended Facility in both the 3033 and 3042 during multiprocessing operations. VM/370 with the VM/System Extensions program product installed supports operation of MVS/System Extensions during multiprocessing operations.

SECTION 05: 3033 PROCESSOR COMPLEX PHYSICAL COMPONENTS AND TECHNOLOGY

PHYSICAL COMPONENTS

The major physical components of a 3033 Processor Complex are shown in Figure 05.1. This uniprocessor configuration consists of a (1) 3033 Processor that contains monolithic processor storage and the channel groups, (2) standalone 3036 Console, and (3) standalone 3037 Power and Coolant Distribution Unit. A static converter or motor generator set is required to supply power for the components of the 3033 Processor Complex. A motor generator set is recommended.

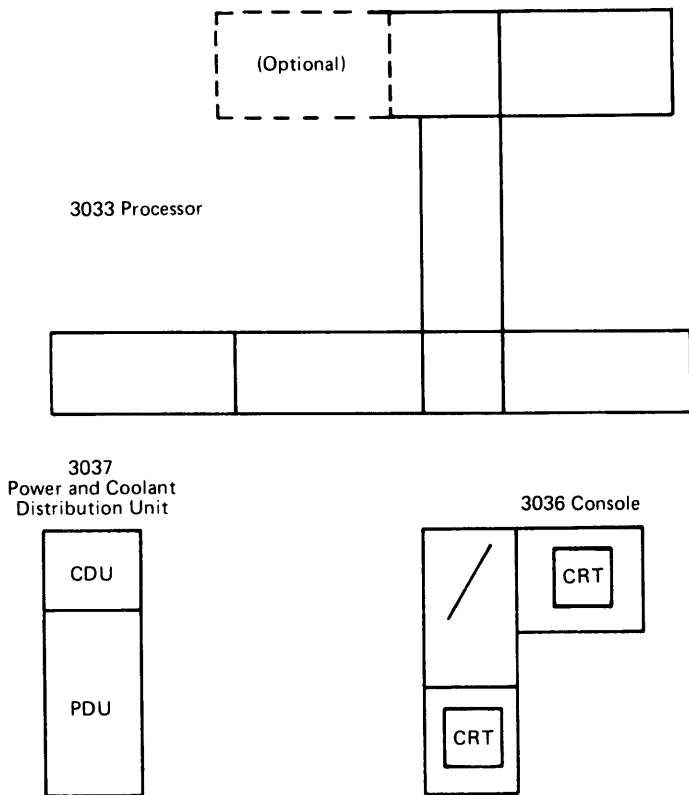


Figure 05.1. Major components of the 3033 Processor Complex

The size of uniprocessor models of the 3033 Processor is the same regardless of the amount of processor storage installed. A 3033 Processor that contains twelve channels does not contain the frame drawn with dotted lines in Figure 05.1. The additional frame is required when the optional third channel group is installed. The size of the 3033 Processor without the optional channel group and director installed is only one frame larger than the 3168 Processing Unit, which does not contain any channels.

Space requirements for the standalone 3036 Console and 3037 Power and Coolant Distribution Unit for a 3033 Processor are approximately the

same as for the 3066 System Console and 3067 Power and Coolant Distribution Unit for the 3168 Processing Unit. Thus, space requirements for a 3033 Processor Complex are significantly less than those for a Model 168 configuration with its standalone channels.

MOTOR GENERATOR SET

A motor generator (MG) set is the converter unit that provides the power required by the components of the 3033 Processor Complex. It takes 60 Hz (cycle) power from the building electrical distribution system, converts it to 415 Hz power, and supplies it to the PDU, from which it is distributed to the other components under control of the 3036 Console.

The MG set should be ordered at the same time as the 3033 Processor Complex, with delivery up to two months prior to installation. While IBM does not manufacture MG sets, a procedure is established for ordering the required MG set through IBM. The same types of MG sets that are used for a Model 168 system and 3032 Processor Complex can be used for a 3033 Processor Complex. Note however, that the 3033 Processor requires approximately 33% less power than the 3168 Processing Unit. (See IBM System/370 Installation Manual - Physical Planning, GC22-7004, for more details concerning MG set size and installation requirements.)

COOLING

The heat generated by the logic boards in the 3033 Processor is removed by forced air and a closed-loop water circulation system. A liquid coolant is used in addition to air because of the amount of heat generated by the densely packed circuits in the 3033 Processor. The combined air and water cooling requirement for a 3033 Processor is approximately 20% less than for the 3168 Processing Unit.

The installation must supply 28 gallons of cooled water per minute (45° to 60° F.) to the coolant distribution unit (CDU), which is housed in the standalone 3037 Power and Coolant Distribution Unit 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 3033 Processor. That is, the user-supplied water does not enter the closed-loop system of the 3033 Processor (see Figure 05.2). The CDU houses the necessary controls to maintain the proper temperature in the closed-loop system. An automatic valve adjusts the flow of chilled water supplied to the CDU.

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 3033 Processor are separated easily, and the valved connections close off automatically to prevent water from escaping.

Physical planning for 3033 Processor installation should ensure that arrangements are made to provide the required water. (See IBM System/370 Installation Manual - Physical Planning, GC22-7004-5, for more details.) Note that the water cooling system used and physical requirements for installation-supplied cooled water for the 3033 Processor are the same as those for a 3032 Processor and 3168 Processing

Unit except for a slight difference in the amount of cooled water that must be supplied.




Figure 05.2. Conceptual flow of the water cooling system in the 3033 Processor

TECHNOLOGY

Monolithic technology is used to implement nearly all logic and all storage (processor, local, reloadable control, and buffer) in the 3033 Processor. Certain components of the 3033 Processor are implemented in a type of technology not used in the Model 168 Model 3 or other System/370 processors. Specifically, technology that is not used in other System/370 processors is used for all logic in the instruction processor function (but not channel groups) in the 3033 Processor. This logic technology is faster and three times more dense than the MST used for logic in the Model 168 Model 3. It accounts for the faster cycle time of the 3033 Processor and also results in reduced space requirements.

While the 3033 Processor has approximately twice the number of circuits as the 3168-3 Processing Unit, a 3033 Processor with twelve channels requires only one frame more than a 3168-3 Processing Unit, which does not contain any channels.

Different technology from that used in the Model 168 Model 3 is also used for the high-speed buffer and control storage of the 3033 Processor. This technology is faster and denser than that used in the Model 168 Model 3 high-speed buffer and control storage and enables the 64K buffer in the 3033 Processor to be implemented in the same amount of space as the 32K buffer for the Model 168 Model 3. In addition, the TLB, high-speed buffer index array, and trace buffer in the 3033 Processor are implemented in a technology that is the same density but three times faster than that used in the same components in the Model 168 Model 3.

The storage chip used for processor storage in 3033 Processors with 12M- or 16M-bytes contains twice the number of bits, 4K instead of 2K, as the storage chip used in most other System/370 processors and in 3033 Processors with 6M-bytes. (A 3033 Processor with 4M- or 8M-bytes may contain 2K or 4K chips in its processor storage, depending on its shipment date.) This double density enables a 3033 Processor to contain 16M-bytes of processor storage in the same amount of space as would be required for 8M-bytes using the 2K-bit chip.

SECTION 10: 3033 PROCESSOR COMPLEX ARCHITECTURE DESIGN AND FUNCTIONAL COMPONENTS

10:05 ARCHITECTURE DESIGN AND PROCESSOR ELEMENTS

ARCHITECTURE DESIGN

Both basic control (BC) mode and extended control (EC) mode are standard in the 3033 Processor. Thus, like other System/370 processors, the 3033 Processor is upward compatible with System/360 when BC mode is used. The 3033 Processor is also upward compatible with other System/370 processors for both BC and EC mode operations.

As a result of the architecture design of the 3033 Processor, BC or EC mode programs written for other System/370 processors will run without modification on a 3033 Processor with a comparable hardware configuration, with the following exceptions:

1. Time-dependent programs. (They may or may not execute correctly.)
2. Programs that depend on the validity of storage data after system power has been turned off and then on.
3. Programs that use processor-dependent data such as that which is logged in the processor-dependent logout area.
4. Programs that depend on the nonusable lower processor storage area being smaller than 1928 bytes. This area can be reduced to 512 bytes by moving the 1416-byte machine check extended logout area.
5. Programs deliberately written to cause certain program checks.
6. Programs that depend on devices or facilities not implemented in the 3033 Processor.
7. Programs that use processor-dependent operations of the 3033 Processor that are not necessarily compatible with the same operations on other System/370 processors.
8. Programs operating in a virtual storage environment using READ DIRECT or WRITE DIRECT instructions that specify a virtual storage address that does not translate to the same real storage address. These programs must execute in virtual=real mode or be modified to specify a virtual=real address in READ DIRECT and WRITE DIRECT instructions. This restriction exists because in a processor with the System/370 Extended Facility installed, the addresses in READ DIRECT and WRITE DIRECT instructions are assumed to be real instead of virtual as in processors without the System/370 Extended Facility installed.

Programs written for System/360 can execute in the 3033 Processor operating in BC mode subject to the same restrictions as System/370 programs. In addition, System/360 programs that use the ASCII mode bit cannot execute in the 3033 Processor. If an OS MFT or MVT control program that was generated for a System/360 processor is used on a 3033 Processor, the processor should be set to check stop on machine checks. When a machine check occurs, the processor stops before a logout to processor storage or the diskette drive occurs.

User-written processing programs that operate in a System/360 model or Model 168 under OS MFT or MVT control and do not violate a restriction can operate in a 3033 Processor under OS/VS1 or OS/VS2 SVS control with little or no modification, as discussed in the optional programming system supplements (Sections 90 and 100). Similarly, user-written processing programs operating under OS MVT can operate under OS/VS2 MVS with little or no modification.

User-written processing programs that operate under one of the System/370 operating systems that support the 3033 Processor and do not violate a restriction can operate on the 3033 Processor without modification using a release of the operating system that supports the 3033 Processor.

PROCESSOR ELEMENTS AND FEATURES

Processor Elements

The major elements in the 3033 Processor and their interrelationships are shown in Figure 10.05.1. Interaction of processor elements with the standalone console is also shown. The major elements are:

- The instruction preprocessing function and execution function, which form the instruction processor function that executes the instruction set for the 3033 Processor
- Processor storage
- The processor storage control function, which controls all access to processor storage by the other processor elements and the console and performs virtual-to-real address translation
- Two or three channel groups and their directors
- Maintenance and retry function

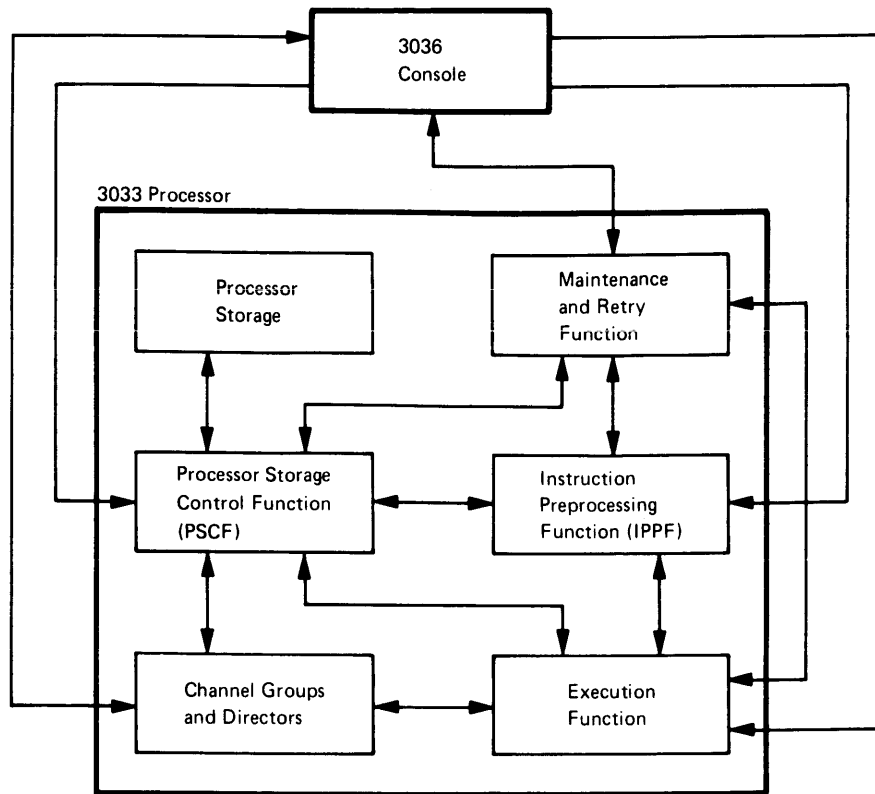


Figure 10.05.1. Major elements in the 3033 Processor

Processor Features

The standard and optional features for the 3033 Processor are listed below. The items not generally available for other System/370 processors or those whose implementation in the 3033 Processor differs from that in the 3168 Processing Unit are identified by an asterisk and discussed in this publication. Items not identified by an asterisk operate in the 3033 Processor in the same manner as in other System/370 processors (see A Guide to the IBM System/370 Model 168 for System/360 Users (GC20-1787) or other appropriate guide series publications).

Standard features for the 3033 Processor are:

- BC and EC mode of operation and control registers
- Instruction set that includes binary, decimal, floating-point, and extended precision floating-point arithmetic. Standard System/370 instructions for the 3033 Processor not provided for System/360 are:

```

CLEAR CHANNEL*
CLEAR I/O
COMPARE AND SWAP
COMPARE DOUBLE AND SWAP
COMPARE LOGICAL CHARACTERS UNDER MASK
COMPARE LOGICAL LONG
INSERT CHARACTERS UNDER MASK
INSERT PSW KEY
INVALIDATE PAGE TABLE ENTRY*

```

LOAD CONTROL
 LOAD REAL ADDRESS
 MONITOR CALL
 MOVE LONG
 PURGE TLB
 RESET REFERENCE BIT
 SET CLOCK
 SET CLOCK COMPARATOR
 SET CPU TIMER
 SET PSW KEY FROM ADDRESS
 SHIFT AND ROUND DECIMAL
 START I/O FAST RELEASE
 STORE CHANNEL ID
 STORE CHARACTERS UNDER MASK
 STORE CLOCK
 STORE CLOCK COMPARATOR
 STORE CONTROL
 STORE CPU ID
 STORE CPU TIMER
 STORE THEN AND SYSTEM MASK
 STORE THEN OR SYSTEM MASK
 TEST PROTECTION*
 MVS-dependent instructions*

- Dynamic Address Translation (includes translator, TLB*, and STO-stack*)
- Reference and Change Recording
- System/370 Extended Facility*
- Instruction retry*
- Interval timer (3.3 ms resolution)
- Time-of-day clock
- Clock comparator and CPU timer
- Monitoring feature
- Program Event Recording
- Program interruption for SSM instruction
- Expanded machine check interruption class*
- ECC on processor storage*
- Byte-oriented operands
- Store and fetch protection
- High-speed buffer storage - 64K bytes*
- Two channel groups, each with six channels (one byte and five block multiplexer)*
- Channel Indirect Data Addressing
- Limited channel logout area with I/O retry data and an extended channel logout*
- Reloadable control storage for the execution function and channel groups
- Store status function
- Direct Control

Optional features for the 3033 Processor, which can be field installed, are:

- Channel-to-Channel Adapter (one in the first group and one in the second channel group)
- Extended Channels (third channel group and director with channels 12 through 15)*
- Two-Byte Interface (for block multiplexer channels 1, 7, and 12 or 13 only)

System/370 Extended Facility

The System/370 Extended Facility provides several extensions to System/370 architecture: (1) low address protection, (2) the TEST PROTECTION (TSPT) instruction, (3) the INVALIDATE PAGE TABLE ENTRY

(IPTE) instruction, (4) the common segment facility, (5) MVS-dependent instructions, and (6) virtual machine extended facility assist.

The IPTE instruction and common segment facility are designed to increase the efficiency of dynamic address translation. They are discussed in Section 10:15.

Low-Address Protection. The low address protection facility is enabled and disabled using bit 3 in control register 0. When enabled, this facility prevents all instructions from storing using addresses in the range of 0 through 511. Checking for the 0 through 511 range is performed before dynamic address translation is performed (if DAT is enabled). A protection program interruption occurs if a store attempt is made in the 0 through 511 address range and execution of the instruction is terminated.

Low-address protection can be used to provide protection against the destruction of data in low processor storage by programs that operate with storage protect key zero. Thus, failures that occur because of inadvertent modifications to locations 0 to 511 are prevented.

Low address protection is applied to the store accesses made by all instructions to addresses 0 to 511 plus the store accesses to the operands of INVALIDATE PAGE TABLE ENTRY and READ DIRECT instructions. Note that when enabled, low-address protection is not applied to store accesses made by the instruction processor or a channel during the following: interruption processing, interval timer updating, machine check and channel logouts, CSW storing during an I/O interruption or after I/O instruction execution, STORE CHANNEL ID instruction processing, IPL processing, or store status processing. It also is not applied to the data stores made by a channel during read I/O operations.

TEST PROTECTION Instruction. The TEST PROTECTION privileged instruction provides a programmed means of testing for the protection exceptions that will occur if a given real storage location is referenced using the specified storage protect Key. The condition code is set to indicate that both fetching from and storing in the location are permitted, only fetching is permitted, or neither fetching nor storing operations are permitted.

This instruction is designed to improve processor performance by eliminating the initialization and termination procedures normally required before a control program tests for potential protection violations before performing service for a processing program.

The TEST PROTECTION instruction specifies the storage location to be tested (a virtual or real address) and the four-bit storage protect key to be used to determine whether storing is allowed. The store protect key in the instruction is compared with the protect key associated with the specified storage location. If they match, storing is permitted unless disallowed by the low address protection facility (that is, low address protection is enabled and the specified address is less than 512, before translation if it is a virtual address). The fetch protect bit in the key associated with the specified address is inspected to determine whether fetching is permitted.

MVS-Dependent Instructions. Twelve privileged instructions that depend upon the conventions, fields, and control block formats of MVS for proper execution are provided to increase MVS performance. These instructions are the following:

- OBTAIN LOCAL LOCK
- RELEASE LOCAL LOCK
- OBTAIN CMS LOCK
- RELEASE CMS LOCK

- TRACE SVC INTERRUPTION
- TRACE PROGRAM INTERRUPTION
- TRACE INITIAL SRB DISPATCH
- TRACE I/O INTERRUPTION
- TRACE TASK DISPATCH
- TRACE SVC RETURN
- FIX PAGE
- SVC ASSIST

Virtual Machine Extended Facility Assist. This assist is designed to improve the performance of virtual machines in which MVS is being used by enabling the System/370 Extended Facility to be utilized in the virtual machine also. This assist enables the twelve MVS-dependent instructions to be executed directly by a virtual machine without an interruption and simulation by CP. This assist is enabled when bit 1 in control register 6 is zero and bit 29 in this register is one. The virtual machine assist and virtual machine extended facility assist can be enabled simultaneously.

10:10 INSTRUCTION PREPROCESSING AND EXECUTION FUNCTIONS

The 3033 Processor has a 57-nanosecond cycle time and an eight-byte-wide data path. Extensive parity checking is performed in the 3033 Processor to ensure the validity of the data being used. All data transfer, logical, and arithmetic operations are parity checked. Automatic hardware retry of most failing instruction operations, without programming assistance, is provided as an availability feature.

INSTRUCTION PREPROCESSING FUNCTION

The 3033 Processor contains an instruction preprocessing function and an execution function that overlap instruction fetching and preparation with instruction execution to improve performance. The instruction preprocessing function is totally controlled by logic circuits and, therefore, can process several instructions concurrently while the execution function is executing a single instruction. Imprecise interruptions do not occur in a 3033 Processor.

The instruction preprocessing function (IPPF) prefetches instructions (maintaining them in sequence), decodes instructions, calculates addresses, prefetches instruction operands, makes estimates of the success of conditional branches, and maintains three instruction streams.

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 preprocessing function. This ensures the availability of prefetched instructions whether the branch is taken or not.

The IPPF contains three sets of instruction buffers each with four doublewords to hold prefetched non-decoded instructions, one instruction register to hold one instruction during its decoding, four instruction queuing registers to hold four decoded instructions until they are transferred to the execution function, a 24-bit adder that can accept three inputs and perform address calculation while decoding occurs, six operand address registers to retain calculated operand addresses for the execution unit, three instruction address registers (one for each instruction stream), an address incrementer for incrementing and decrementing addresses, and a length incrementer for computing the end-address of storage operands. Assuming an average instruction size of four bytes, 24 instructions can be prefetched and held in the IPPF, of which up to four can be decoded.

While the design of the instruction preprocessing function in the 3033 Processor is similar to that used in the instruction unit of the Model 168 processor, the instruction preprocessing function in the 3033 Processor contains the following major enhancements that provide increased internal performance over the Model 168:

- Three instruction streams of four doublewords are maintained instead of two instruction streams of two doublewords. This allows up to two branch instructions (instead of one as in the Model 168) to be processed concurrently.
- A copy of general registers 1 to 15 (contained in local storage) is maintained in the instruction preprocessing function to use for address calculations that are performed by the instruction preprocessing function. This enables instruction decoding and operand address generation for an instruction to be done in one cycle in the 3033 Processor while two are required in the Model 168.
- The 3033 Processor has six operand address registers and six corresponding doubleword buffers, instead of two of each as in the Model 168. This substantially improves the ability of the IPPF to prefetch operands.
- The processing of CVD, CLC, MC, LM, STM, STNSM, STOSM, ICM, CLCM, STCM, STCTL, STIDP, SCK, SCKC, STCKC, SPT, STPT, STPX, and STAP instructions is overlapped with the processing of subsequent instructions in the 3033 Processor. Overlap of the following is disabled for only two cycles unless a word-overlap condition exists: MVN, MVC, MVZ, NC, OC, and XC. Overlap is disabled for all these instructions in the Model 168.
- The execution of store type instructions is speeded up by the addition of certain registers. In the Model 168, there are two F-registers (for data storage) and each is associated with a specific operand address register. Only one register set can be used for an instruction. In the 3033 Processor, there are four F-registers and six operand address registers. There is no fixed relationship between a given F-register and address register and an instruction can utilize more than one register set. Thus, for example, an MVC instruction can store every other cycle and its execution improves by a factor of two to one. Execution of an STM instruction improves by a factor of four to one.
- The IPPF obtains data not contained in the buffer faster than the instruction unit in a Model 168. In the 3033 Processor, as each of the eight doublewords required to fill a buffer block are sent to the buffer, the IPPF can obtain the required doubleword(s) after the first without waiting for the entire buffer block to be loaded and then fetching the required doubleword(s), as is done in the Model 168. This particularly improves instruction fetching and the execution of an instruction that accesses more than one doubleword (LM, for example).
- Significantly less performance degradation occurs in the 3033 Processor when byte-oriented operands are utilized because of changes in the way these operands are aligned (IPPF does aligning instead of a microprogrammed routine in the execution function, for example).
- Under certain conditions, the decoding of an instruction following a LOAD or INSERT CHARACTERS UNDER MASK (with a mask of 7 or F) can be performed one cycle earlier in the 3033 Processor than in the Model 168 processor. The cycle can be saved when a decoded, unexecuted L or ICM instruction alters the general register that is specified as

the index or base register in a successive instruction about to be decoded.

- Execution time for various instructions is reduced by one or more cycles through various other setup changes in the IPPF.

The instruction preprocessing function in the 3033 Processor also contains an instruction pretest function like that in the instruction unit in a Model 168 for use in instruction nullification operations when dynamic address translation is operative. In the 3033 Processor, pretesting is performed by IPPF hardware and/or additional microcode routines that are executed before normal instruction execution. However, for some instructions, prefetching of data accomplishes pretesting so that no additional pretesting cycles are required. A LOAD instruction that addresses a word on a fullword boundary is an example of such an instruction.

EXECUTION FUNCTION

The execution function is predominantly microprogram controlled and can execute one instruction at a time. It has the capability of processing a new instruction every processor cycle and many instructions (approximately half) can be executed in one cycle. Emphasis is placed on optimizing 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, as in the execution unit of the 3168 Processing Unit.

The execution function in the 3033 Processor differs from the execution unit in the 3168 Processing Unit in the following major areas:

- High-speed multiply, an optional feature for the Model 168 that provides faster execution of binary and floating-point arithmetic operations, is basic to the implementation of the execution function.
- Storage-to-storage instruction overlap and general register manipulation is speeded up by the use of six doubleword operand buffers, instead of two as in the Model 168 execution unit.
- More instructions require only a single execution cycle than in the Model 168.
- The processing of certain storage-to-storage format instructions, such as CLC, NC, OC, and XC, when the possibility of destructive overlap does not exist, is faster. In the 3033, the IPPF aligns operands and the execution function uses the parallel adder to process eight bytes at a time (in two cycles). In the Model 168, processing is done one byte per cycle.
- MVCL and CLCL instructions execute faster by processing eight bytes per cycle, instead of eight bytes in four cycles. This enables more data to be processed before execution of the instruction must be stopped to take an interruption and thereby reduces the amount of instruction stop and start-up processing.
- Execution time for ED and EDMK instructions is significantly improved through the implementation of a new algorithm that requires more control storage.
- Execution time for other instructions, such as AP, SP, CP, and MVO is reduced through microprogram changes that require additional control storage.

- Local storage is modified to permit an odd/even or even/odd general register pair to be read at a time, instead of one register at a time.
- Execution time for various other instructions is reduced by one or more cycles as a result of the improved setup implemented in the IPPF and/or changes in processing the instruction in the execution function.

Local storage contains the general and floating-point registers and has a read/write cycle time of 57 nanoseconds. It can be accessed by four sources and written into from one source in the same processor cycle.

Control storage for the execution function consists of 4K words of monolithic reloadable control storage (RCS), with a 57-nanosecond cycle time. RCS can contain the entire instruction set for the 3033 Processor.

During a processor power-on or a processor initial microprogram load (IMPL) sequence, execution function RCS is automatically loaded with microcode from a removable diskette contained on a diskette drive in the 3036 Console. See Section 20:10 for a further discussion of microprogram loading.

Instruction Retry

Detected hardware errors in the processor, except those that occur during execution of certain instructions that have passed beyond a threshold point, can be retried automatically by the instruction retry capability of the execution function. A mask bit (9) in control register 14 (which controls the logging of retry data to the fixed logout area) determines whether the processor is enabled or disabled for the instruction retry function.

If the processor is enabled for instruction retry, retry occurs after instruction errors, after failures that occur during interruption 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 interruption if the instruction has not proceeded beyond an established threshold point.

Instruction 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.

Instruction retry is accomplished by additional microprogram routines in the execution function and hardware included in the maintenance and retry function. The failing processor 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 microprogram. However, certain SS-format instructions are retried from the point of successful execution. In this case, the microprogram saves the status data necessary to restart at the proper byte.

The following instructions cannot be retried in the 3033 Processor: DIAGNOSE, READ DIRECT, WRITE DIRECT, LOAD CONTROL, SET STORAGE KEY, RESET REFERENCE BIT, TEST AND SET, SET PREFIX, SIGNAL PROCESSOR, INVALIDATE PAGE TABLE ENTRY, the subject instruction of an EXECUTE, and

an INSERT CHARACTERS UNDER MASK instruction that modifies the register used in calculating its operand address.

Two instruction retry mode control bits (which are set via a DIAGNOSE instruction) and a system recovery mask bit (4) in control register 14 determine whether a machine check interruption occurs after a successful instruction retry when the instruction retry facility is enabled. The retry mode control bits operate subject to the setting of the system recovery mask bit. That is, the system recovery mask bit must be on as well when the mode control bits are set to recording mode in order for interruptions to occur after successful retries.

Interruptions

As in other System/370 processors, the implementation of the machine check class of interruption in the 3033 Processor is expanded considerably from its implementation in System/360 in order to enhance system availability. The other four interruption classes (I/O, SVC, program, and external) operate in the same manner in the 3033 Processor as in System/360 processors except for (1) additional program and external interruptions that are defined for new features, (2) expansion of channel masking, and (3) expansion of external interruption masking to support new features.

When all interruption classes are enabled, the high-to-low priority for honoring requests is exigent (hard) machine check, program or SVC, repressible (soft) machine check, external, and I/O.

Machine Check Interruptions. The 3033 Processor presents one of seven subclasses of machine check interruption, depending on the specific processor malfunction. Each interruption subclass is maskable and causes either a repressible machine check or an exigent machine check interruption (formerly called soft or hard machine check interruption, respectively) when enabled and a logout to the fixed processor storage area.

Fixed processor storage in the 3033 Processor consists of the fixed locations in decimal addresses 0-191, the fixed logout area in locations 256-351, the machine check interruption extended information in locations 216-255 and 352-511, and the machine check extended logout area of 1416 bytes, which normally begins at location 512. Fixed locations 0-127 are identical in layout and content to these locations in other System/370 processors. Figure 10.10.1 shows the processor-dependent logout area for the 3033 Processor.

A logout to the fixed locations, machine check interruption extended information area, and (if bit 9 in control register 14 is on) processor-dependent fixed logout area occurs when any type of machine check interruption is taken. The data is processed by recovery management routines. The machine check interruption extended information area data indicates the reason for the interruption in the machine check code and contains the contents of the general, floating-point, and control registers as well as CPU timer and clock comparator values.

The processor-dependent machine check extended logout area begins at the address specified in control register 15, which is set to decimal location 512 during IPL or system reset. The save areas in the machine check extended logout area preserve the status of the processor at the time of the machine check interruption.

The occurrence of a logout to the machine check extended logout area is controlled by the setting of the extended logout mask bit (8) in control register 14. Fixed logouts must also be enabled via bit 9 in control register 14 in order for an extended logout to occur. When

enabled, an extended logout occurs before the first and after the seventh instruction retry or when an exigent machine check interruption takes place. This data can be recorded by recovery management routines, and the Processor Logout Analysis program can be used to process this data at a later time.

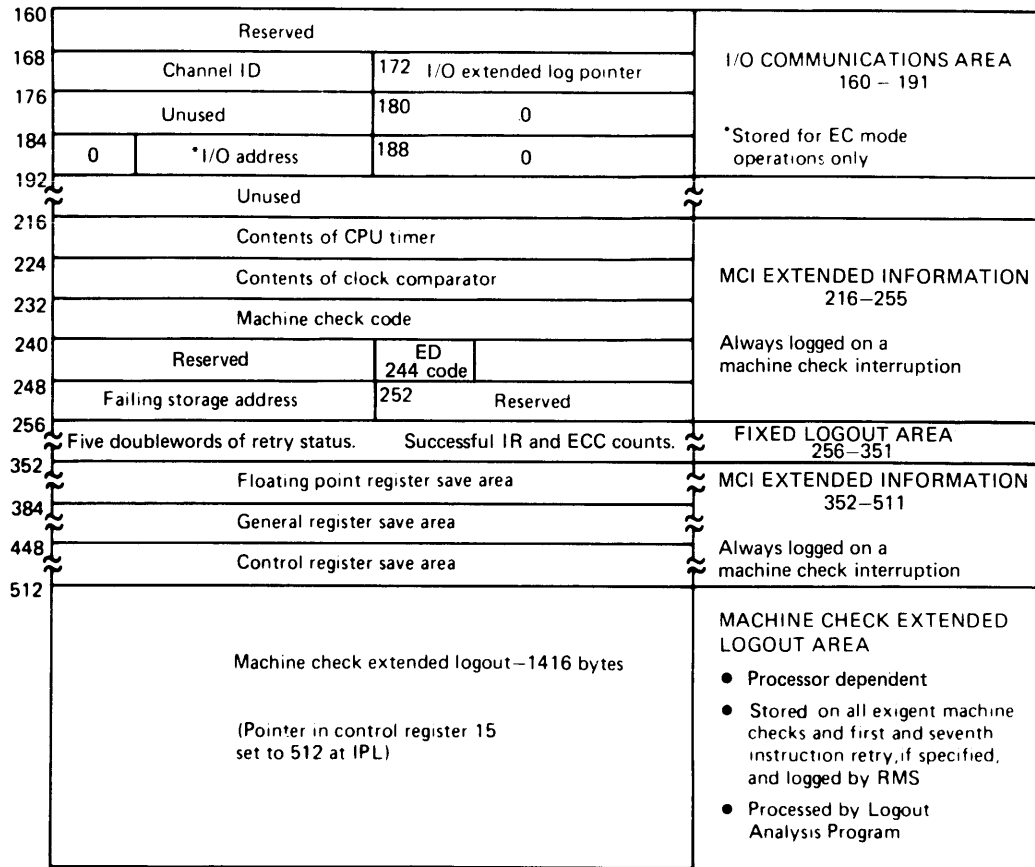



Figure 10.10.1. Processor-dependent fixed storage locations

Figure 10.10.2 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 interruption occurred, the validity of certain fields stored in the fixed logout area, and the length of the machine check extended logout area.

Table 10.10.1 lists the machine check subclasses defined for the 3033 Processor. They are described in the discussion that follows. The mask bits used to enable or disable the processor for interruptions for each subclass are indicated, and the setting of the machine check code is discussed. PSW bit 13 and three other mask bits are used to enable and disable the processor for machine check interruptions. The recovery mask (R), external damage mask (E), and degradation (D) mask bits are contained in control register 14 and operate subject to PSW bit 13. If PSW bit 13 is off, the processor is disabled for all machine checks. If PSW bit 13 is on, the settings of the three additional mask bits determine whether or not interruptions, other than system damage and instruction processing damage, will be taken.

Fixed Logout Area Locations 232-239



4

Bit	Interrupt Type	Bit	Error	Bit	Valid Fixed Area Data
0	SD - System Damage	14	Backed Up	20-23	Machine Check Old PSW (48-55)
1	PD - Instruction Processing Damage	15	Delayed interruption		20 AMWP
2	SR - System Recovery	16	Storage Error		21 Masks and Protect Key
3	TD - Timer Damage		Uncorrected		22 Program Mask and Condition Code
4	CD - Timing Facilities Damage	17	Storage Error	24	23 Instruction Address
5	ED - External Damage		Corrected	27	Failing Storage Address
7	DG - Degradation	18	Protection Key Error	28	Floating Point Registers (352-383)
				29	General Registers (384-447)
				30	Control Registers (448-511)
				31	Machine Check Extended Logout Storage (Validity of storage being processed by instructions when interruption occurred.)
				46	CPU Timer Value
				47	Clock Comparator Value

Figure 10.10.2. 3033 Processor machine check code

Repressible machine check interruptions for the 3033 Processor are as follows:

- **System Recovery.** This interruption can occur if both PSW bit 13 and the recovery mask bit (4 in control register 14) are on. It is caused by a successful instruction 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 instruction retry of a failing instruction processor operation (including one caused by a buffer malfunction). The SC bit (bit 17) will be on as well and the failing storage address is stored in locations 248 to 251 if an ECC recovery occurred. Bit 24 in the machine check code indicates the validity of the stored address. Error recording and, possibly, buffer deletion are required after a system recovery interruption.

- **Timer Damage.** This interruption can occur if PSW bit 13 and the external damage mask bit (6 in control register 14) are on. It indicates damage to the interval timer. Programmed validation procedures and error logging are required.
- **External Damage.** This interruption can occur if both PSW bit 13 and the external damage mask bit (6 in control register 14) are on. There are four conditions that cause this interruption. The specific reason for the error is indicated in the external damage

code that is stored in location 244. The external damage (ED) bit (5) and the external damage validity bit (26) in the stored machine check code are on after an external damage machine check interruption.

One condition that causes an external damage condition is an uncorrectable error in the PSCF or processor storage (such as a double-bit error) that occurred during an I/O operation and that did not affect instruction processor function execution. A one in bit 2 in the external damage code indicates this error condition.

Another external damage interruption condition is a hard error in a channel group. This condition exists when one or more channels enter the not-operational state without performing an I/O system reset on the interface. This error prevents any channels in the group from continuing operation and is indicated by a one in bit 3 of the external damage code.

The limited channel logout and an I/O extended logout (if enabled) are also stored when a channel group error occurs. The limited channel logout contains a validity bit (15) that indicates whether the stored I/O extended logout is valid. Programmed recovery procedures can be attempted after a channel group error using the CLEAR CHANNEL instruction (see the discussion in Section 10:20 under "I/O Error Handling".)

The other two conditions that cause an external damage interruption are the elapsing of the timeout interval associated with the execution of an I/O instruction or I/O interruption. If an I/O instruction is not executed in the eight-second timeout interval or an I/O interruption does not occur in the eight-second timeout interval, bit 5 or bit 6, respectively, in the external damage code is on.

- Timing Facilities Damage. This interruption can occur if both PSW bit 13 and the external damage mask bit (6 in control register 14) are on to indicate an error condition in the time-of-day clock, CPU timer, or clock comparator. The CD bit in the stored machine check code (bit 4) will be on to indicate a timing hardware error.

This interruption is generated when an error occurs 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 a clock failure.

When a STORE CLOCK COMPARATOR or STORE CPU TIMER instruction is issued and the addressed timing facility has an error, or when the CPU timer or clock comparator develops an error, this interruption is taken.

- Degradation. This interruption can occur if PSW bit 13 and the degradation mask bit (7 in control register 14) are on. The DG bit (7) will be on in the stored machine check code to indicate that a buffer row or half the TLB was deleted.

Table 10.10.1. 3033 Processor machine check interruptions

Subclass	Mask Bit(s)	Cause	Machine Check Condition
System Recovery	PSW 13 and R (and instruction retry or ECC mode bits set to recording mode)	<ul style="list-style-type: none"> • Instruction processor error corrected by retry • Single-bit processor storage error corrected by ECC 	Repressible
Timer Damage	PSW 13 and E	<ul style="list-style-type: none"> • Interval timer damage 	Repressible
Timing Facilities Damage	PSW 13 and E	<ul style="list-style-type: none"> • Error in time-of-day clock, CPU timer, or clock comparator 	Repressible
External Damage	PSW 13 and E	<ul style="list-style-type: none"> • Error that did not affect the instruction processor (such as a double-bit processor storage error during an I/O operation) or a channel group error 	Repressible
Degradation	PSW 13 and D	<ul style="list-style-type: none"> • An error occurred that resulted in the deletion of a buffer row or half of the translation lookaside buffer 	Repressible
Instruction Processing Damage	PSW 13	<ul style="list-style-type: none"> • One of the following and none of the conditions listed for system damage: Unretryable processor error Uncorrectable processor error Double-bit processor storage error Storage protect key failure 	Exigent
System Damage	PSW 13	<ul style="list-style-type: none"> • One of the instruction processing damage conditions <u>and</u> one of the following: LPSW instruction in execution Interruption in progress Hang-detect condition 	Exigent

Exigent machine check interruptions for the 3033 Processor are as follows:

- **Instruction Processing Damage.** This interruption can occur if PSW bit 13 is on. The PD bit (1) in the stored machine check interruption code is used to indicate an error occurred during the execution of the instruction indicated by the machine check old PSW. The error was either a double-bit processor storage failure, a storage protection failure, an instruction processor error that was unretriable, or an instruction processor error that could not be corrected by the instruction retry hardware, with none of the conditions that cause a system damage interruption present also.

If a double-bit processor storage failure caused the interruption, the SE bit (16) in the stored machine check code is on also, and the address of the failing storage area is indicated in locations 248 to 251. Bit 24 in the machine check code indicates the validity of the stored address. A storage protection failure is indicated by the KE bit (18).

If enabled, a logout to the machine check extended logout area also occurs. For an instruction processor error, an extended logout occurs before the first and after the seventh retry. Error logging and the execution of recovery procedures are required after this type of interruption.

- **System Damage.** This interruption can occur if PSW bit 13 is on. The SD bit (0) is on in the stored machine check code to indicate an error occurred during the execution of the instruction indicated by the machine check old PSW. This interruption is indicated instead of instruction processing damage when an LPSW instruction is in execution, an interruption is in progress, or a hang-detect condition exists in addition to one of the conditions that causes an instruction processing damage interruption. Logouts occur as for instruction processing damage and recovery procedures are required.

Two modes of system operation for machine check interruptions are possible: full recording mode and quiet, or nonrecording, mode. In full recording mode, the processor is enabled for all machine check interruption types and all cause an interruption to be taken and logouts to occur. In quiet mode, the processor is disabled for certain machine check interruptions for repressible conditions. Quiet mode can be used to permit system operation without machine check error recording for all or certain repressible conditions when a large number of transient (correctable) errors are occurring. It can also be used to allow 3033 Processor operation under the control of an operating system without 3033 Processor machine check handling routines included.

A check-stop state and associated check stop control bit (0 in control register 14) are defined for the 3033 Processor. If PSW bit 13 is zero and the check stop bit is one when a system damage condition occurs or if a system damage condition occurs during the processing of a previous machine check error, processor operations cease immediately without the occurrence of a logout to the fixed or extended logout processor storage area. Implementation of a check stop prevents system operations from continuing when the nature of the processor malfunction prevents the processor from presenting meaningful status data.

The state of the 3033 Processor after IPL or a system reset is:

1. The processor is disabled for instruction retry (bit 9 in control register 14 is 0). An exigent machine check condition results on any instruction processor error, including those caused by a buffer malfunction. The instruction retry mode control bits

indicate recording mode after an IPL and the system recovery mask is off.

2. The processor is disabled for system recovery interruptions (system recovery mask bit in control register 14 is 0). Therefore, single-bit processor storage error corrections do not cause a repressible machine check condition. Note that the ECC mode control bits are set to indicate nonrecording mode after an IPL so that even if the system recovery mask bit is turned on, an interruption will not occur after successful single-bit error corrections unless the ECC mode control bits are changed.
3. The processor is enabled for external damage interruptions (external damage mask bit is 1). A PSCF or processor storage failure associated with an I/O operation, a hard channel group error, an I/O instruction or I/O interruption timeout, or damage to the interval timer, time-of-day clock, CPU timer, or clock comparator causes a machine check interruption.
4. The degradation mask bit is off. Buffer row deletions and a TLB half deletion will not cause a machine check interruption.
5. PSW bit 13 is normally set on by the IPL PSW (it is set off by system reset) so that system damage machine checks (unretryable instruction processor failures, unsuccessfully retried instruction processor errors, storage protection failures, and double-bit processor storage errors associated with the instruction processor) cause an exigent machine check interruption.
6. The check-stop control bit is on.
7. Fixed logouts and machine check extended logouts are disabled (bits 8 and 9 in control register 14 are 0) and control register 15 points to location 512 as the beginning of the machine check extended logout area.
8. Extended channel logouts are not enabled.

10:15 PROCESSOR STORAGE AND THE PROCESSOR STORAGE CONTROL FUNCTION

PROCESSOR (MAIN) STORAGE

The 3033 Processor has a two-level storage system--a high-speed buffer storage backed by a large processor (main) storage. The 3033 Processor has eight-way, doubleword interleaved processor storage with cycle times of 285 and 456 nanoseconds buffered by 64K of high-speed monolithic buffer storage with a 57-nanosecond cycle time.

Processor storage sizes of 4M-, 8M-, 12M-, and 16M-bytes (Models U4, U8, U12, and U16), where M equals 1,048,576, are provided for the 3033 Processor in a uniprocessor configuration. (Model U6 is no longer available.) Model changes are field installable. In order to achieve interleaving, the processor storage present in the 3033 Processor is divided into logical storage elements. A logical storage element is generally defined as that portion of processor storage that can operate independently from all other processor storage.

The data path to and from processor storage is eight bytes wide. Eight consecutively addressed doublewords are spread across logical storage elements 0 through 7 so that 64 consecutive bytes can be accessed concurrently (one doubleword from each of the eight logical storage elements, requested one 57-nanosecond cycle apart). That is,

processor storage location 0 is in logical storage element 0, location 8 is in logical storage element 1, and location 56 is in logical storage element 7. Processor storage location 64 is in logical storage element 0, and the address distribution sequence continues through all available storage locations.

Processor storage can be accessed concurrently by any combination of one or more channels and the instruction processor function for a total of eight unique logical storage requests. Contention arises when two or more components attempt to access the same logical storage element simultaneously. In addition, requests made to a busy logical storage element are delayed until the storage element becomes free.

Every processor cycle, processor storage priority is established within the processor storage control function for concurrent requests to each non-busy logical storage element. The channel bus controller, which controls all requests for processor storage from all the channels in the processor, has the highest priority for a non-busy logical storage element.

An instruction processor function fetch request has the next highest priority after the channel bus controller. Instruction processor function store requests (four can be outstanding concurrently) have the lowest priority and the stores are made in first-in, first-out sequence. However, if the instruction processor function fetch request is for the same storage location as any one of the outstanding instruction processor function store requests, the instruction processor function fetch request is given lower priority than the instruction processor function store requests.

The storage distribution element (SDE) directs the flow of data between the processor storage control function and the logical storage elements. It also provides the ECC logic used for the contents of processor storage and controls store and fetch protect checking.

The storage protect array is contained in the processor storage element. It contains one seven-bit key for each 2K of processor storage (four store-protect bits, one fetch-protect bit, one reference bit, and one change bit). The operation of store and fetch protection and reference and change recording is the same in the 3033 Processor as in other System/370 processors with DAT hardware.

As in a 3168 Processing Unit, protection checking is accomplished in the 3033 Processor using the protect key in the translation lookaside buffer (TLB) if the translation of a virtual storage address is found in the TLB. Otherwise, the storage protect array is used.

Error checking and correction (ECC) hardware, which is contained in the processor storage element, provides automatic detection and correction of all single-bit processor storage errors and only detection of all double-bit and many multiple-bit errors. The ECC feature operates in the 3033 Processor in the same manner as in other System/370 processors; however, a new ECC code is implemented in the 3033 Processor. This code reduces the potential number of undetected multiple-bit errors by 50% over the number undetected in 3168 processor storage.

The system recovery mask bit in control register 14 and two ECC mode control bits, which are set using the DIAGNOSE instruction, control 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 interruption does not occur after the successful correction of a single-bit error. The mode indicated in the ECC mode control bits is subject to the setting of the system recovery mask bit. Full recording

mode is effective only when the system recovery mask bit is also on when full recording mode is set in the ECC mode control bits.

Table 10.15.1 summarizes the cycle and access times for the 3033 Processor. The cycle time indicates how frequently a given logical storage element can be accessed. The fetch times indicate the time required to make data from processor storage available in a register in the instruction processor function where it can be utilized. Fetch and store times assume the required logical storage element is not busy at the time of the request.

Table 10.15.1. 3033 Processor cycle and access times

Cycle or Access Time	Time in Nanoseconds
Processor cycle time	57
Local storage cycle time	57
Control storage cycle time	57
Processor (logical) storage read/write cycle time (for eight bytes on a doubleword boundary)	285
Processor storage cycle time for a partial write (fewer than eight bytes)	456
Minimum time between successive selects to processor storage	57
Processor storage access time (from time of PSCF select to availability of data in the PSCF)	285
Instruction processor function fetch of eight bytes from processor storage (from time of request acceptance to availability of data in an instruction processor register) assuming the logical storage is not busy	456
Instruction processor function fetch of eight bytes from buffer (from time of request acceptance to availability of data in an instruction processor register)	114
Minimum time between successive buffer requests	57

Processor Storage Reconfiguration

The processor storage present in the 3033 Processor is divided into from two to eight bands of 2M bytes each. Band numbers 0 through 7 are used to identify the bands. The first 2M-byte band is 0, the second is 1, etc.

The configuration frame of the 3036 Console is used to enable storage bands, assign a two-megabyte range of addresses to each enabled band, and establish eight-way interleaving or serial operations. Serial operation of processor storage is designed to be used primarily by customer engineers. These processor storage configuration changes can be made only at system reset time.

Any 2M-byte address range can be assigned to any enabled band. The address range 0 to 2M must be assigned and the operating systems that support the 3033 Processor Complex require consecutive address ranges to be assigned to enabled bands.

If an uncorrectable processor storage error occurs, the operator has the ability to remove the malfunctioning storage band from the operative processor storage and reconfigure the addresses of the remaining bands to achieve consecutive storage addressing. Then the operating system can be re-IPLed and the system can continue operating with reduced available processor storage.

PROCESSOR STORAGE CONTROL FUNCTION

The processor storage control function (PSCF) controls all access to processor storage. In general, the IPPF makes fetch requests for instructions and their operands while the execution function makes data store requests, although the execution function can also make operand fetch requests under certain conditions. The console and maintenance and retry function make both fetch and store requests.

The components of the PSCF and its interconnection with other processor elements and the standalone console are shown in Figure 10.15.1. The PSCF contains the high-speed buffer and its controls, dynamic address translation hardware (translator, translation lookaside buffer, and segment table origin address stack), and the channel bus controller. Priority and contention controls for processor storage and processor storage configuration logic are also contained in the PSCF. The function of the channel bus controller is discussed in Section 10:20, which describes the channels in the 3033 Processor.

High-Speed Buffer Storage

The high internal performance of the 3033 Processor is achieved partly by the inclusion of a 64K-byte high-speed buffer storage unit. The buffer provides high-speed data access for instruction processor function fetches. The higher internal performance of the 3033 Processor versus the 3168-3 Processing Unit is partly the result of the use of a 64K, instead of a 32K, high-speed buffer.

The buffer has a 57-nanosecond cycle. The instruction processor function can obtain eight bytes from the buffer in two cycles, or 114 nanoseconds, and a request can be initiated every cycle. This is the time between request acceptance and availability of the data in an instruction processor register. This compares with a time of 456 nanoseconds required to fetch eight bytes of data directly from processor storage if the required data is not currently in the buffer.

Buffer storage control and use are handled entirely by hardware and are 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 3033 Processor (a least recently used algorithm) is like that used in the 3168 Processing Unit buffer.

The processor storage control function contains the high-speed buffer and controls all buffer and processor storage references made by the instruction processor function, channels, and console. The buffer storage control portion of the PSCF handles instruction processor function-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 instruction processor function, buffer storage control determines whether or not the requested data is in the high-speed buffer by interrogating its address array of the contents of the buffer. If the data requested is present in the buffer, it is sent directly to the instruction processor function 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 instruction processor function. The data is also assigned a buffer location and stored in the buffer. When data is stored by the instruction processor function, 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.

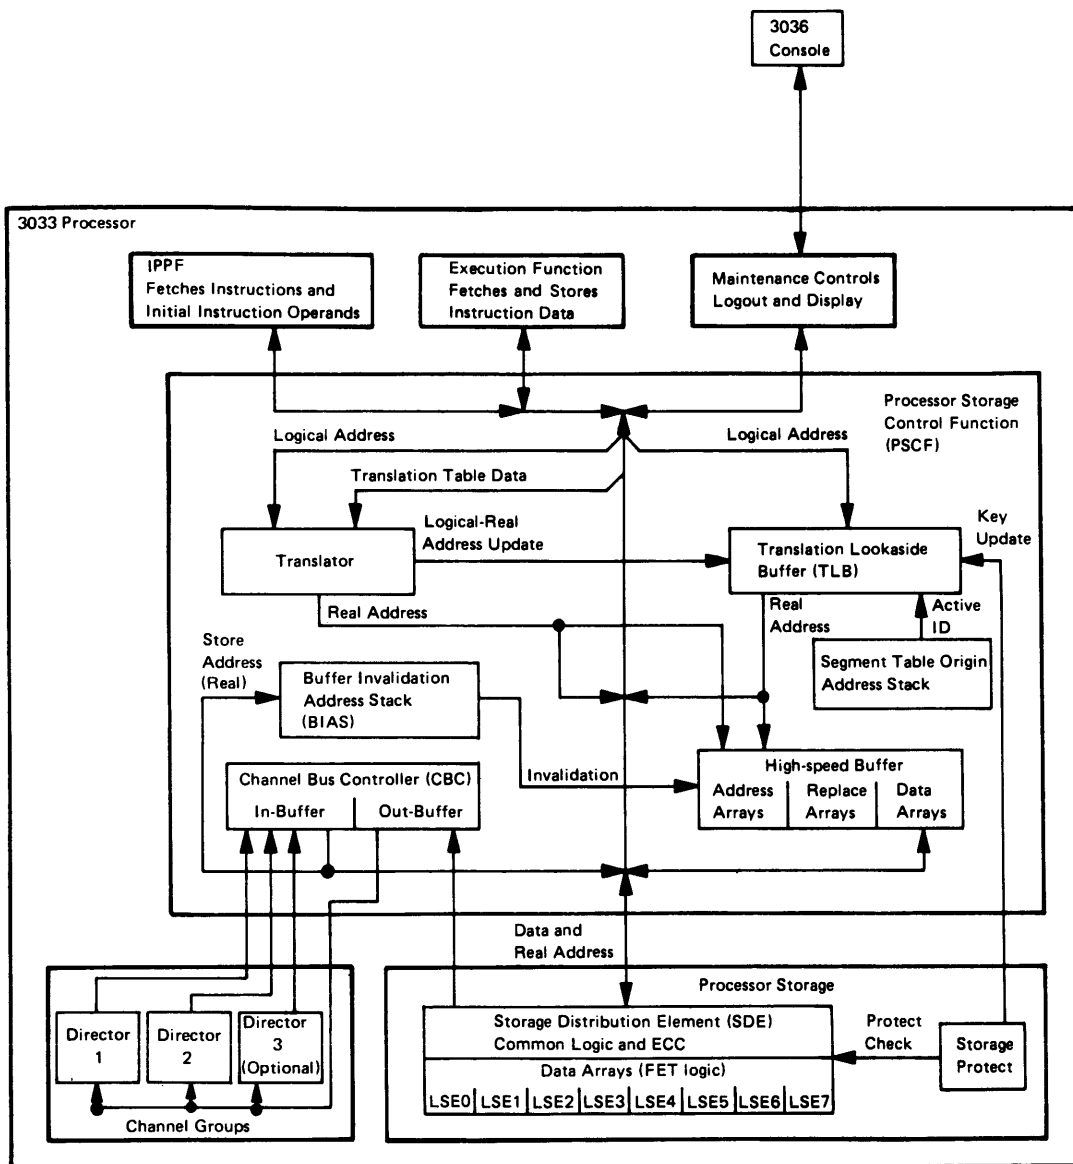


Figure 10.15.1. Processor storage control function components and interface with other processor elements

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, appropriate bits are set in the address array to indicate this buffer data is no longer valid.

The entire buffer can be disabled by the operator using the 3036 Console. When the buffer is disabled, all fetches are made directly from processor storage and effective instruction execution speed is reduced. In addition, a bit can be set by a DIAGNOSE instruction that causes the entire buffer to be bypassed during processor storage fetches and stores without resetting its contents. Selective disabling of a

portion of the buffer when a buffer error occurs (buffer block and buffer row deletion) is also implemented.

The 64K buffer is shown in Figure 10.15.2. The buffer contains 64 columns, each of which is subdivided into 16 blocks. A block is 64 bytes (eight doublewords) and can contain 64 consecutive bytes from processor storage that are on a 64-byte boundary. The buffer can contain a maximum of 1024 different blocks of processor storage data (16 blocks per column times 64 columns). There are 16 rows in the buffer. The first row consists of block 0 of each column (64 blocks). The last row consists of block 15 of each column.

A valid bit is associated with each buffer block and is set to indicate whether or not the block contains valid data. All valid bits are set off during a buffer reset.

Processor storage is logically divided into the same number of columns as buffer storage (64). While there are 16 blocks in a buffer column, the number of blocks in a processor storage column varies with the size of processor storage. When buffer storage is assigned, bits 20-25 of the processor storage address determine which one of the 64 columns in buffer storage is to be used.

The organization of processor storage is shown in Figure 10.15.2. Any of the blocks in a given processor storage column can be placed in any one of the 16 blocks in a corresponding buffer column. Figure 10.15.3 shows the format used for buffer storage addressing.

Buffer contents and buffer block assignment are controlled by an address array, as shown in Figure 10.15.2, and a replacement array. The address array for the buffer is divided into 64 columns consisting of 16 block address registers each, such that there is 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. It also contains a block valid bit and block delete bit. When an instruction processor function-to-processor storage reference is made, the 16 appropriate address array column registers (13-bit block addresses) are interrogated to determine whether the requested data is currently in the buffer and valid.

A replacement array is used to maintain knowledge of the activity of the data blocks within each of the buffer columns. The array consists of 64 logic-controlled activity lists, one list for each column in the buffer. An activity list contains bits that are used in a modified least-recently-used algorithm that is different from the least-recently-used buffer replacement algorithm implemented in the Model 168.

An activity list is altered whenever the instruction processor function references one of the buffer blocks associated with the list. When a block within a given column in the buffer has to be assigned and loaded because data requested by the instruction processor function is not in the buffer, the bits in the activity list for the column are decoded to determine the least recently used block. Thus, the more active data is maintained in the buffer.

The buffer operates as follows. When the instruction processor function requests data, bits 20-25 of the processor storage address of the data are used to obtain a buffer column address. The 13 high-order bits of the processor storage address are then compared with the address in each of the 16 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 bit for that block is on,

the appropriate doubleword from the buffer block is sent to the instruction processor function as determined by bits 26-28 of the processor storage address. A processor storage reference is not made. The entries in the appropriate column activity list in the replacement array are resequenced to reflect this reference.

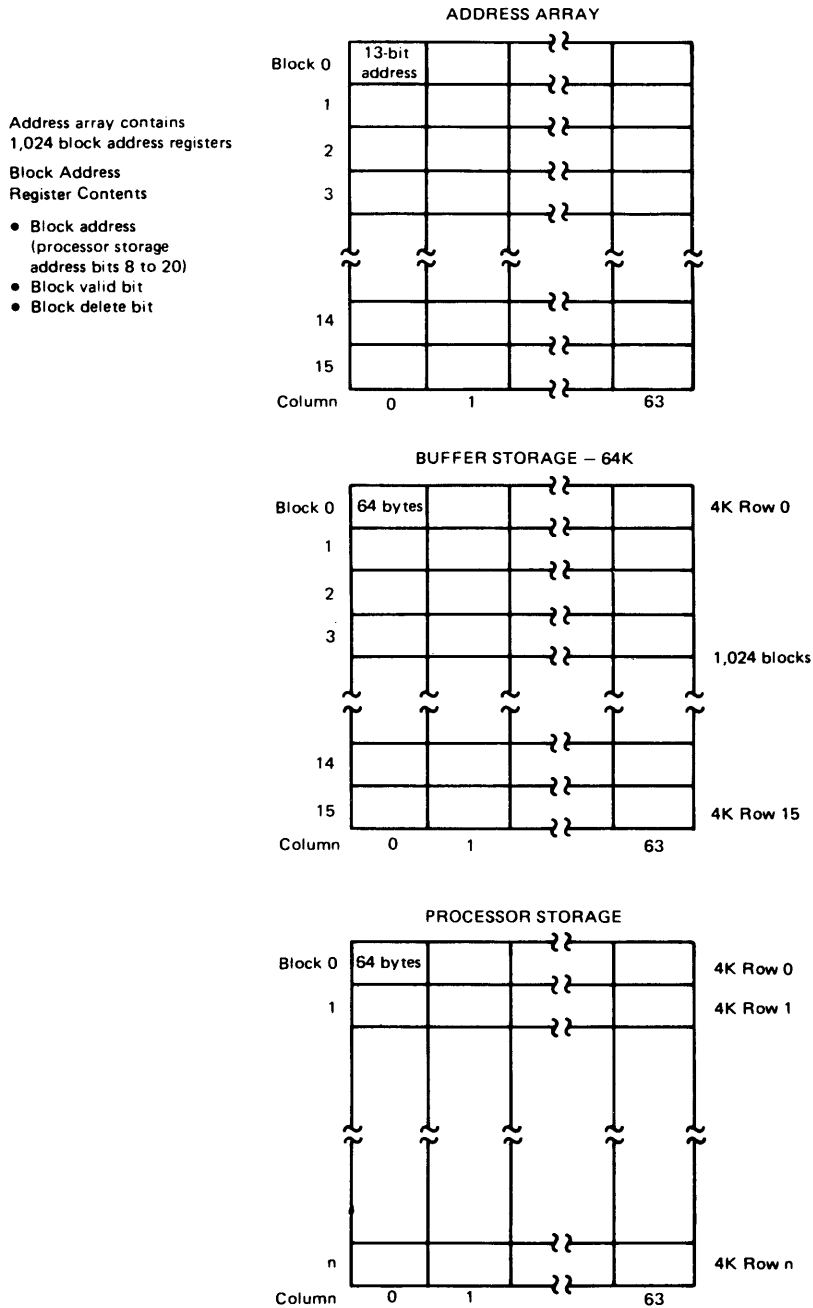


Figure 10.15.2. 64K buffer organization

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

8-20	Used for address compare
20-25	Used to reference buffer column
26-28	Used to reference doubleword within a block
29-31	Used to reference byte within a doubleword

Figure 10.15.3. Processor storage address format for buffer reference

If the desired data block is not in the buffer column interrogated, the requested data must be fetched from processor storage, sent to the instruction processor function, and stored in the buffer. The replacement array activity list for the column involved is decoded and the buffer block determined to be the least recently used is assigned to receive the requested data from processor storage.

Eight processor storage references, one cycle apart, are made to obtain the eight consecutive doublewords and place them in the assigned buffer block. The valid bit 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 instruction processor function. It is sent to the instruction processor function as well as to the buffer so that processing can continue as soon as possible.

Two buffer deletion facilities are implemented in the 3033 Processor: programmed buffer block deletion and automatic buffer row deletion. A delete bit is associated with each block in the buffer. These delete bits are in the address array. When a buffer block fails, it can be deleted by the execution of a DIAGNOSE instruction that turns on its delete bit. When the delete bit for a buffer block is on, the block is no longer selected to receive data when a block must be assigned. All block delete bits are turned off during IPL or system reset.

Automatic buffer row deletion is controlled by a bit that is set via the DIAGNOSE instruction. After IPL or system reset, the setting of this bit enables buffer row deletion. When the bit is set to disable automatic buffer row deletions, any rows currently disabled are reenabled.

Each row in the buffer also has a row delete bit associated with it. When certain errors occur in a row, the row delete bit is automatically turned on when the row deletion facility is enabled. When the delete bit is on for a row, data can no longer be stored in or fetched from any block in the row. This facility permits only 4K of the buffer to be deleted when certain malfunctions occur instead of bypassing the entire buffer.

The 64K buffer operates at its 64K capacity when the 3033 Processor is operating with dynamic address translation mode disabled or with dynamic address translation mode and a 4K page size enabled. When dynamic address translation and a 2K page size are enabled, the buffer operates at a 32K capacity (only columns 0 to 31 are used).

The reason for using a 32K capacity when a 2K page size is enabled is the following. Bits 20 to 25 of the referenced processor (real) storage address are required to determine the column address (0 to 63). When a 4K page size is used, bits 20 to 31 of the referenced virtual storage address are the same as bits 20 to 31 of the corresponding real storage

address and do not need to be translated. However, when a 2K page size is used, bit 20 must be translated as only bits 21 to 31 in the virtual and corresponding real storage addresses are equal.

Therefore, if the 64K capacity were to be used for a 2K page size, bit 20 would not be available for buffer address array column addressing until after address translation had been performed. By using a 32K capacity for a 2K page size, bits 20 to 25 are available for accessing the appropriate buffer address array column before address translation is performed.

Whenever the buffer in the 3033 Processor is reset, it is set to operate at its 64K capacity. The buffer is reset when one of the following occurs: IPL, program reset, power on reset, reset at the end of a machine check interruption, system reset, or system clear. When a LOAD CONTROL instruction is issued to change the page size in effect from 2K to 4K, the buffer is also reset. When page size is changed from 4K to 2K, a reset is not performed but buffer capacity is reduced to 32K.

Dynamic Address Translation

The dynamic address translation (DAT) facility of the 3033 Processor supports the same virtual storage size (16M-byte maximum) and organization (2K or 4K pages and 64K or 1024K segments) as other System/370 processors with DAT hardware. The translation process using segment and page tables also operates as in other System/370 processors and the same addresses are translated. The only change in the segment and page tables used in the 3033 Processor is the use of bit 30 in each segment table entry to support the common segment facility.

Reference and change recording using the store and fetch protect key are performed as in other System/370 processors and instruction nullification is accomplished as in the 3168 Processing Unit.

As in certain other System/370 processors, a translation lookaside buffer (TLB) and segment table origin address stack (STO-stack) are implemented in the 3033 Processor to speed up the address translation process. Operation of these two facilities can be enabled and disabled via programming using the DIAGNOSE instruction using a service frame of the 3036 Console. Operation of the STO-stack can be enabled and disabled separately from the TLB when the TLB is enabled. The IPTE instruction and common segment facility are also implemented to increase address translation efficiency.

Translation Lookaside Buffer. The translation lookaside buffer is used to retain up to 128 previously translated virtual storage addresses. Addresses associated with up to 29 different virtual storages can be contained in the TLB at any time.

The TLB for the 3033 Processor is shown in Figure 10.15.4. It contains an A and B array, each of which contains 64 locations addressed 0 to 63. A pair of locations, one in array A and the correspondingly addressed location in array B, is always addressed at one time. Each location in each array can contain a virtual storage address and its translated real storage address. Each array can contain translations for virtual storages for which a 2K page size is being used and for virtual storages for which 4K is the page size.

Every time a virtual storage address is translated during instruction execution, the virtual storage address, the resulting real storage address and its associated storage protect key, and identification of the virtual storage to which the virtual storage address belongs (see STO-stack discussion) are placed in one of the 128 TLB locations.

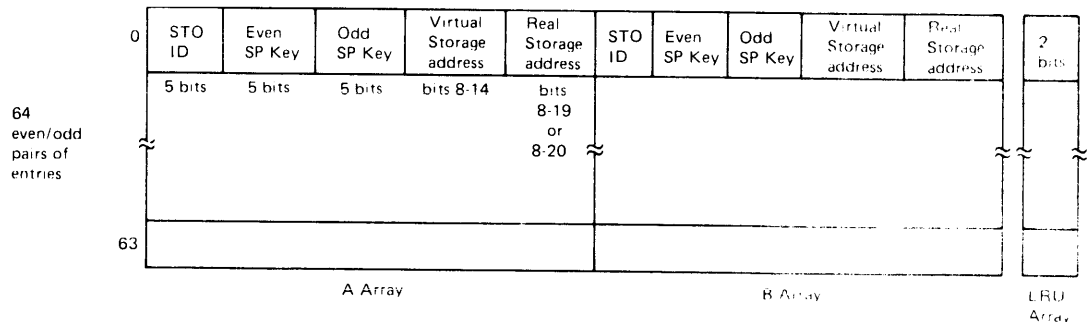


Figure 10.15.4. TLB format

A hashing algorithm is applied to bits 9 through 20 of the virtual storage address to obtain a six-bit address that selects one of the 64 TLB locations (A and B array location pair). A 64-location least recently used (LRU) array is associated with the TLB that determines whether the A or B array location is assigned. Locations 0 to 63 in the LRU array correspond to and are addressed a half-cycle later than the 64 A and B array location pairs. Each LRU array location contains two bits that identify the least recently used translation of the two in the corresponding TLB location pair.

Each time a translation is referenced in the TLB, the corresponding LRU array bits are updated. When a translation is performed, after the TLB location pair is determined the array assigned is the one indicated by the LRU array bits, which reflects the translation used longest ago for the selected TLB location pair. Thus, the TLB will contain the most recently referenced translations.

The TLB in the 3033 Processor is oriented towards a 4K page size. Thus, only one of the 128 TLB locations is used for a translation regardless of the page size in effect. In the Model 168, a translation for a virtual storage that is using a 4K page size is placed in two TLB locations.

After the effective virtual storage address has been computed and before performing the translation using segment and page tables, the TLB is interrogated to determine whether it contains the required translated address. Interrogation of the TLB is done in parallel with reference to the index array for the buffer, as previously indicated. Therefore, no translation cycles are required when the translated address is obtained from the TLB. If the TLB does not contain the required translation or if the entry is invalid, as indicated by a zero identification code, the complete table-lookup translation procedure, using segment and page tables, is performed.

In the 3033 Processor, the number of processor cycles required for address translation when the translation is not obtained from the TLB varies from a minimum of 10 to a maximum of 40, assuming no I/O interference, depending on the locations (buffer storage or processor storage) of the segment table and page table entries required for the translation.

If an error occurs in the TLB, the array containing the error (64 locations) is disabled and a machine check interruption occurs if degradation interruptions are enabled. The degradation bit will be on in the stored machine check code. The disabled half of the TLB is reenabled during an IPL, system reset, or when a bit that controls

whether or not the TLB disabling function is active is set to disable the function. The control bit is set via a DIAGNOSE instruction and is the same bit that controls automatic high-speed buffer row deletions. On a system reset, the control bit is set to enable the TLB delete function.

When a SET STORAGE KEY instruction is issued and valid translated addresses are in the TLB, the entire TLB is searched (by hardware) and each entry is invalidated (set to all zeros) that has the same real address as the one for which the key is being set. Whenever such invalidation occurs, the LRU array entries for the TLB location pair containing the invalidated entry is set to indicate that the array containing the invalidated entry is the least recently used.

Two instructions are provided to perform programmed TLB invalidation. When the INVALIDATE PAGE TABLE ENTRY (IPTE) instruction is issued, the specified page table entry is invalidated and the TLB is inspected by hardware for entries that used the now invalid page table entry. The LRU array is updated as appropriate for each entry found. This instruction eliminates the need to purge the entire TLB when only one page table entry is invalidated.

The PURGE TLB instruction is provided to enable a program to invalidate all 128 TLB entries. Purging the TLB consists of writing zeroes in all 128 locations of the TLB, setting all entries in the LRU array to indicate array A as the least recently used, and setting a TLB indicator to specify the TLB does not contain any translated addresses. This TLB indicator is tested when an IPTE or SSK instruction is issued, to avoid searching of the TLB when it does not contain any translated addresses.

STO-Stack. A change in segment table origin address, segment size, or page size can also affect the validity of current TLB entries. In order to reduce the number of full TLB purges required by such changes, a segment table origin address register stack (STO-stack) is implemented. The common segment facility provides the ability to use TLB entries more efficiently under certain conditions.

The STO-stack can contain the addresses of 29 different segment tables at a time. Each segment table could define a different virtual storage. An STO-stack entry also indicates the segment and page size in effect for the virtual storage associated with the segment table address.

The 29 entries in the STO-stack have a unique identification number, 2 to 30, associated with each. One of these numbers (the one associated with the currently active virtual storage) is denoted to be the currently active identification number. Whenever a segment table address is placed in control register 1, the segment table address is also placed in the STO-stack, if it is not already there, and the identification number assigned to the segment table address becomes the new active identification number.

An STO-stack identification number is stored with each TLB entry to identify the segment table, and thereby the virtual storage with which the TLB entry is associated. The identification number 31 is used to identify an entry that is associated with a common segment.

When the common segment facility is used, bit 30 in each segment table entry indicates whether the associated segment is private (bit 30 is 0) or common (bit 30 is 1). The common segment bit must be set by the routine that maintains the segment table. A segment can be identified as common when its pages have the same page frames assigned in every virtual storage in which it is contained and, thus, its virtual

storage addresses always translate to the same real storage address regardless of the set of segment and page tables used.

When a translation is placed in the TLB, bit 30 in the segment table entry used in the translation process is inspected. If it is on, the STO-stack identification stored in the TLB entry is 31 to indicate the entry belongs to a common segment. An identification of 31 indicates the TLB entry can be utilized regardless of the currently active identification number because it contains a translation that is valid for each virtual storage being used. When bit 30 in the segment table entry is zero, the identification of the virtual storage to which the entry belongs (currently active identification) is stored in the TLB entry.

When the TLB is interrogated to determine whether it contains the required translation, the STO-stack identification number of each TLB entry is compared with the active identification number. If the identifications are equal, the TLB location contains a translation from the virtual storage associated with the active identification number and a comparison is made between the virtual storage address to be translated and that in the TLB entry. When the identification in the TLB location is 31, an equal comparison is forced and a virtual storage address comparison is made.

If the identifications are not equal, the TLB location contains a translation for a virtual storage other than the one associated with the currently active identification number and, therefore, the TLB entry does not contain the required translation even though it may contain a virtual storage address equal to the one that is to be translated. When an equal comparison occurs for both the identification number and the virtual storage address in a TLB entry, the translation is taken from the TLB; otherwise, the translation procedure that uses segment and page tables is performed.

When (1) DAT mode is entered or (2) when DAT mode is operative and a LOAD CONTROL instruction is issued for control register 1 or to load bits 8 to 12 of control register 0, the segment table address in control register 1 and page and segment size specifications from control register 0 are compared with each of the STO-stack locations to determine whether a change in these specifications is being made.

An equal comparison between any valid STO-stack entry and the segment table address, segment size, and page size in control registers 0 and 1 indicates that the virtual storage associated with the segment table address now in control register 1 is currently one of the virtual storages whose translations are being maintained in the TLB and that segment and page size have not been changed. The STO-stack identification number of the segment table address now in control register 1 is designated to be the active identification. No TLB purging is required.

No equal comparison between any valid STO-stack entry and the segment table address, segment size, and page size in control registers 0 and 1 indicates that translations for the segment table now indicated by control register 1 are not currently being maintained in the TLB or that segment or page size for a virtual storage whose segment table address is in the STO-stack is being changed. The segment table address, segment size, and page size in control register 1 are placed in the STO-stack, and the STO-stack identification number assigned becomes the active identification.

STO-stack locations are assigned on a consecutive basis beginning with the location assigned to identification number 2. A next available STO-stack location register is updated each time a segment address is placed in the next available STO-stack location. When all 29 STO-stack

locations are full and a segment address must be added, a reset procedure is performed. The entire TLB is purged, the new segment address, segment size, and page size are placed in the location for identification number 2, the register indicating the next available STO-stack location is set to 3, and the active identification number is set to 2.

Note that the STO-stack entries are never physically invalidated. The next available STO-stack location register value indicates which STO-stack entries are valid. When this register is set to a value of 2, as after a STO-stack reset operation, this indicates no STO-stack entries are valid. A value of 6, for example, indicates the STO-stack entries for identification numbers 2 through 5 are valid, etc.

An identification number of 2 is assigned to all TLB entries while the STO-stack is disabled. When the STO-stack is reenabled, the reset procedure performed when the STO-stack is full (described above) is executed.

All entries in the TLB are purged and the STO-stack is reset whenever one of the following occurs:

- System reset or CPU reset (even if the TLB and STO-stack are disabled)
- The PURGE TLB instruction is issued or a machine check interruption occurs (even if the TLB is disabled)
- SET PREFIX instruction is issued
- The STO-stack is disabled or enabled
- The TLB and STO-stack are enabled
- The STO-stack is full and a new segment address is to be added to the stack

Implementation of the STO-stack in the 3033 Processor enables a control program that supports multiple virtual storages (such as OS/VS2 MVS and VM/370) to alter control registers 0 and 1 in order to change the virtual storage for which address translation is effective, without automatically causing purging of the entire TLB. The STO-stack facility will also be of benefit in an OS/VS2 SVS environment, since OS/VS2 SVS supports two segment tables to provide fetch protection for all regions.

The common segment facility can be utilized to further reduce TLB purging when multiple virtual storages are supported in the manner used by MVS, for example. Certain fixed MVS control program areas that appear in the same virtual storage locations in each virtual storage supported can be identified as common segments when they are allocated real storage on a virtual address equals real address basis.

10:20 CHANNELS

CHANNELS AND DIRECTORS

The channels in the 3033 Processor are packaged in groups of six and four. A channel group is controlled by a director, which is a microprogram-controlled processor that provides common logic and control functions for the channels. A director has its own reloadable control storage, arithmetic logic unit, and storage areas. A channel group can be serviced concurrently with normal processing operations and it can be powered on and off separately from other channel groups and processor components for card changing.

Two groups of six channels each and their associated directors are standard in the 3033 Processor. Each standard channel group includes one byte multiplexer channel and five block multiplexer channels. In the first channel group, the byte multiplexer channel is addressed as 0 and the block multiplexer channels as 1 through 5. In the second channel group, the byte multiplexer is addressed as 6 and the block multiplexers as 7 through 11.

Optionally, one additional channel group and director consisting of four block multiplexer channels or one byte and three block multiplexer channels, addressed as 12, 13, 14, and 15, can be added to the 3033 Processor (via the Extended Channels feature) for a total of 16 channels maximum. The channel addresses for the three channel groups are fixed and cannot be altered.

The byte and block multiplexer channels provided in the 3033 Processor are functionally equivalent to the byte and block multiplexer channels for other System/370 processors. The channel indirect data addressing feature, provided for use when dynamic address translation mode is operative, is standard in the 3033 Processor and operates in the same way as in other System/370 processors.

The byte multiplexer channel can operate in byte mode, permitting several slow-speed devices to operate concurrently, or in burst mode, permitting one high-speed device to operate. A byte multiplexer channel is capable of a data transfer rate generally in the range of 40 to 75 KB/sec. Byte multiplexer channel configurations consisting of control units with fast I/O interface turnaround times and multibyte data transfer may achieve a higher range. Those configurations consisting primarily of control units with slow I/O interface turnaround times and single-byte data transfer will achieve a lower range. Individual and aggregate channel data rate capability depends on the I/O device configuration and block multiplexer load within the channel group.

A block multiplexer channel can operate in selector mode to allow the operation of one device at a time or in block multiplexer mode to allow the concurrent operation of multiple channel programs, as indicated by the setting of bit 0 in control register 0.

A block multiplexer channel in a 3033 Processor can operate at a data rate of up to 1.5 MB/sec. When the optional Two-Byte Interface feature is installed on a block multiplexer channel, its maximum data rate is increased to up to 3 MB/sec. The Two-Byte Interface feature enables 3-megabyte data rate devices, such as the 2305 Model 1, to be attached to the 3033 Processor and permits the 3838 Array Processor to transfer data to and from the 3033 Processor at a data rate up to 3 MB/sec.

The Two-Byte Interface feature can be installed on only the first block multiplexer channel in each channel group (channels 1, 7, and 12 or 13). Optionally, one or two Channel-to-Channel Adapters, which are functionally equivalent to those provided for other System/370 processors, can be installed in a 3033 Processor. One adapter can be attached to any channel in the first channel group while the other can be attached to any channel in the second channel group. While the optional channel group can not have a Channel-to-Channel Adapter installed, one of its channels can be connected to a Channel-to-Channel Adapter that is installed in another processor.

Each director controls the channel-to-processor storage communication (via the channel bus controller) and channel-to-device communication (via the I/O interface) for the channels in its group. Each director contains the following:

- Reloadable control storage that contains the microprogram for the director and its channels. The director and each of its channels

operates within its own microprogram in the shared reloadable control storage. The microprogram is shared by the switching of control at specified points in the microprogram. A control switch is the temporary halting of the execution of the current microprogram to give control to another microprogram. Reloadable control storage for each director and its channels is loaded during a processor IMPL sequence.

- An arithmetic logic unit for channel and data control functions.
- Controls and buffers for the channels controlled by the director.
- Local storage that is shared by the director and the channels. One portion of this local storage is used only by the channels and contains a section dedicated to each channel. The other portion of the local storage is shared by the director and the channels in its group. Local storage is used as a working area for the director microprogram, for unit control word (UCW) storage and updating, and to contain channel control information for the I/O operations in progress for the channel group.
- A UCW storage area consisting of 1536 UCWs, for a six-channel group, or 1024 UCWs, for a four-channel group. Each channel in a group has 256 UCWs available to it (one for each of the 256 I/O device addresses possible for a channel).

The block multiplexer channels in a group are controlled by a microprogram that is shared and by special hardware controls. I/O instructions, initial device selection, data transfers to and from processor storage, channel command word (CCW) fetching while command chaining, and I/O interruptions for block multiplexer channels are microprogram controlled.

The special block multiplexer channel hardware controls are used for data buffer-to-I/O interface data transfers and for the command chaining reselection sequence. Thus, I/O interface data handling for one channel is performed without interfering with concurrent microprogrammed operations on any other channel in the group using the special hardware controls.

The byte multiplexer channel is microprogram controlled for all operations except device selection, which is hardware controlled. Byte multiplexer microprogram operation can be interrupted to give control to another microprogram when a block multiplexer channel in the group needs servicing.

Since each director has its own control storage and arithmetic logic unit, channel operations in the 3033 Processor do not interfere with instruction processing operations except in the case of simultaneous requests for access to processor storage elements. A channel request for a given non-busy logical storage element is given priority over an instruction processor request for the same logical storage element.

UCW Assignment

Each director contains enough UCWs to provide 256 UCWs for each channel in its group. The UCWs available are designated as shared or nonshared. The byte and block multiplexer channel UCWs to be shared are specified by the customer engineer using the 3036 Console (director configuration frame).

Shared UCW assignment is first performed during system installation but can be altered by the customer engineer any time thereafter if required by an I/O configuration change. The assignments are written to the operational diskettes that contain microcode for the processor.

For a byte multiplexer channel, all 256 UCWs can be nonshared or up to eight can be designated as shared UCWs. When a byte multiplexer UCW is configured for sharing, the maximum number of devices that can share the UCW (8, 16, or 32) must also be specified. The number of nonshared UCWs available for a byte multiplexer channel with shared UCWs is 256 less the maximum number of devices specified for each shared UCW (that is, less 8, 16, or 32 per shared UCW).

Byte multiplexer shared UCWs always permit disconnection during a command chained channel program regardless of whether block multiplexing mode is enabled for the block multiplexer channels.

The UCWs for the five block multiplexer channels in a channel group also can be all nonshared or some can be designated as shared. Each block multiplexer channel can have up to eight shared UCWs assigned. When sharing is specified for a block multiplexer UCW, the maximum number of devices that can share the UCW (8, 16, or 32) and the mode of operation of the UCW (selector or block multiplexer) must be indicated also. The number of nonshared UCWs available for the five block multiplexer channels in a standard channel group when UCW sharing is specified is 1280 less the number of shared block multiplexer UCWs (up to 40 maximum) assigned.

When all 256 UCWs for a byte or block multiplexer channel are nonshared, device addresses 00 to FF are used and each device attached to the channel is assigned a unique UCW.

When a block (or byte) multiplexer UCW is designated as shared, a contiguous block of addresses must be used by the control unit assigned to the shared UCW. These addresses can have one of the following ranges, depending on the maximum number of devices that can share the UCW:

- X0 through X7 or X8 through XF for 8 devices maximum
- X0 through XF for 16 devices maximum
- X0 through (X+1)F for 32 devices maximum

When a shared block multiplexer UCW is assigned selector mode of operation, there is no disconnection during the operation of a command chained channel program whether the channel is operating in selector or block multiplexer mode. When block multiplexer mode is assigned to a shared UCW, there is disconnection during the execution of a command chained channel program when the channel is operating in block multiplexer mode and the device is capable of disconnecting. Block multiplexer mode for shared block multiplexer UCWs is designed for control units such as the 3272.

A UCW table is maintained to indicate the current UCW assignments. This table contains one entry for each of the UCWs in the processor (3072 for the two standard channel groups and 1024 for the optional channel group when it is installed). Each UCW table entry is associated with an I/O address for a specific channel, points to the location of its associated UCW, indicates whether the UCW is shared or nonshared, and for a block multiplexer shared UCW indicates whether the UCW is to operate in selector or block multiplexer mode.

During system reset, the UCW table is initialized to indicate each UCW is nonshared. Then the UCW table is modified to reflect the shared UCW assignments that were specified by the customer engineer using the director configuration frame of the 3036 Console. Once processing begins, each time a START I/O is issued to a device, the entry in the UCW table for the device is inspected to obtain any needed information.

System performance is affected by the type of UCW, shared or nonshared, and mode (for shared block multiplexer UCWs) allocated to devices attached to a block multiplexer channel. Table 10.20.1 indicates for the most frequently used I/O devices the (1) type of channel to which the device can be attached, (2) preferred type of channel to which the device should be attached, (3) preferred UCW type when attached to a byte multiplexer channel, and (4) preferred UCW type and mode when the device is attached to a block multiplexer channel operating in block multiplexer mode. Where selector is specified in the channel attachment capability column, it means the selector mode of a block multiplexer channel can be used.

Table 10.20.1. Channel attachment and UCW type/mode for the more frequently used I/O devices that attach to the 3033 Processor

I/O Device or Control Unit	Channel Attachment Capability	Recommended Channel Type	Recommended UCW type for the byte multiplexer channel	Recommended UCW type/mode when block multiplexer mode is used for the block multiplexer channel
Card Punches, Printers, and Diskette				
2501/2520	Byte, Sel, Block	Byte	Nonshared	Nonshared
2821/3811	Byte, Sel, Block	Block	Nonshared	Nonshared
3505/3525	Byte, Sel, Block	Block	Nonshared	Nonshared
1442 N1, N2	Byte, Sel, Block	Byte	Nonshared	Nonshared
1443	Byte, Sel, Block	Byte	Nonshared	Nonshared
3540	Byte, Sel, Block	Byte	Nonshared	Nonshared
3800	Byte, Sel, Block	Block	Nonshared	Nonshared
Magnetic Character Readers				
1419	Byte, Block	Byte	Nonshared	Nonshared
3890	Byte, Block	Block	Nonshared	Nonshared
Optical Character Readers				
1287/1288	Byte, Sel, Block	Byte	Nonshared	Nonshared
3886	Byte, Sel, Block	Byte	Nonshared	Nonshared
Displays				
2250/2840	Byte, Sel, Block	Byte	Shared	Shared (selector mode)
2848	Byte, Sel, Block	Byte	Shared	Shared (selector mode)
3272	Byte, Sel, Block	Block	Shared	Shared (block multiplexer mode)
3250	Block	Block	-	Shared (block multiplexer mode)

Table 10.20.1. (continued)

I/O Device or Control Unit	Channel Attachment Capability	Recommended Channel Type	Recommended UCW type for the byte multiplexer channel	Recommended UCW type/mode when block multiplexer mode is used for the block multiplexer channel
Magnetic Tape				
2803	Sel,Block	Block	-	Shared (selector mode)
3803	Sel,Block	Block	-	Shared (selector mode)
Consoles				
3036	Byte,Sel,Block	Byte	Shared	Shared (selector mode)
2150/1052M7	Byte,Sel,Block	Byte	Nonshared	Nonshared
Direct Access Storage				
2314/2319	Sel,Block	Block	-	Shared (selector mode)
2835 Model 1	Block	Block	-	Nonshared
2835 Model 2	Block	Block	-	Nonshared
2841	Sel,Block	Block	-	Shared (selector mode)
3830	Block	Block	-	Nonshared
3851 A models	Block	Block	-	Nonshared
3851 B models	Byte,Block	Byte	Nonshared	Nonshared
Communications				
2701	Byte,Sel,Block	Byte	Nonshared	Nonshared
3704	Byte	Byte	Nonshared	-
3705-Channel Adapter Type 1	Byte	Byte	Nonshared	-
3705-Channel Adapter Type 2, 3, or 4	Byte,Sel,Block	Block	Nonshared	Nonshared
2702,2703,2715	Byte	Byte	Nonshared	-
Audio Response Unit				
7770 Model 3	Byte	Byte	Nonshared	-
Array Processor				
3838	Block	Block	-	Nonshared
Channel-to-Channel Adapter	Byte,Sel,Block	Block	Nonshared	Nonshared

Channel Priorities and Configuring

The priority for handling simultaneous requests for channel operations within a channel group, high to low, is the following:

- Block multiplexer channel data transfer
- Block multiplexer channel data chaining
- Block multiplexer channel command chaining
- Byte multiplexer channel operations
- Director operations

When simultaneous requests for a data transfer, data chaining, or command chaining operation are made by block multiplexer channels in the same group, the requests are serviced in ascending channel number sequence, with channel 1 (channel 7 in the second channel group and channel 12 or 13 in the third channel group) having the highest priority. Hence, if block multiplexer channel 3 requests a data transfer operation simultaneously with command chaining requests from block multiplexer channels 1 and 2, the order of servicing the requests is channel 3, channel 1, channel 2.

Simultaneous I/O interruption requests from channels are handled in ascending channel number sequence, with channel 0 having the highest priority and channel 11 (or 15 if the third channel group is present) having the lowest.

When the channel group configuration rules discussed below are observed, the five block multiplexer channels in each standard channel group can generally achieve a maximum aggregate data rate of 6.7 MB/sec. A maximum aggregate data rate of 5.5 MB/sec can generally be achieved by the four block multiplexer channels in the optional channel group. Thus, for a 3033 Processor with three channel groups the maximum aggregate data rate for the 14 block multiplexer channels is 18.9 MB/sec.

The stated aggregate data rates can generally be achieved with negligible effect on system performance due to channel overrun. For most channel programs and I/O devices, recovery from overrun is handled automatically by channel and control unit hardware. However, in some cases recovery is handled by programming.

In order to achieve the 6.7 MB/sec data rate for a channel group, the following configuration rules should be followed for the group:

- Attach devices to the block multiplexer channels within a channel group in the following sequence, starting with the lowest numbered channel:
 1. Direct access devices in the following order: 2305 Model 2, (attached to a maximum of two channels), 3350, 3330-series, 3340, and 2314 disk storage.
 2. Magnetic tape units. When magnetic tape is attached to more than one block multiplexer channel, the highest data rate tapes should be attached first.
 3. Buffered devices (3211 and 3800 printers, 3505 card reader, a Channel-to-Channel Adapter, the 3705 operating in NCP mode, for example). When buffered devices are attached to the same block multiplexer channel, the devices with the shortest critical time should be attached first and those with the longest critical time last. The critical time for a buffered device is generally defined as the time required for a channel from/to buffer operation (load a line to a printer buffer, unload a card image from a card reader buffer, for example).

- If block multiplexer channels within the group have (1) a mixture of direct access device types attached to the same channel, or (2) a mixture of direct access devices and buffered I/O devices attached to the same channel, these channels should be considered to be direct access device only channels and configured using the sequence listed in item 1 under the first bullet above.
- Channel utilization (activity) should be balanced across all available block multiplexer channels in the processor.
- When the Two-Byte Interface feature is installed on the first channel in a group, the maximum aggregate data rate for the group can be achieved with negligible effect on system performance due to channel overrun when direct access storage is attached to the first channel and a maximum of three other channels in the group. The preceding guidelines must also be observed.

In addition to the listed configuration rules, when data chaining is to be used, other factors must be taken into account. Specifically, chaining checks for READ CCWs and overruns on WRITE CCWs can occur on unbuffered or buffered devices that do not respond to Suppress Data, when data chaining is specified, because of data address alignment, short CCW byte counts, device data rates, or other channel activity within the director.

To reduce the probability of chaining checks and overrun during data chaining operations, the following recommendations apply:

- Preferably, data addresses should be on quadword boundaries; at a minimum they should be on word boundaries.
- The CCW byte counts should be as large as possible; at a minimum they should be equal to, 54 times S where S is equal to the device data rate in megabytes per second.

The above discussion is not intended to preclude normal data chaining operations between the count, key, and data fields of a record on a direct access device. Each installation should review the 3033 channel configuration with its IBM representative.

Note that certain errors related to a channel affect only that channel and operations on other channels in the group continue normally. However, any permanent error that affects a director results in the loss of all channels in the associated group (see discussion under "I/O Error Handling"). Since two channel groups are standard in the 3033 Processor, channel switching features can be installed for critical I/O devices so they can be accessed via the other standard operational channel group (and/or the third optional group) to process a critical subset of the workload.

The microprogrammed director and channel group approach offers certain advantages over the standalone (2860, 2870, and 2880) channels provided for Model 168 processors. First, significantly less floor space is required for the 3033 Processor and its channels than for a Model 168 and its standalone channels. Second, cabling requirements are reduced and the channels in a 3033 Processor have smaller power and cooling requirements than the Model 168 standalone channels. These benefits are obtained without losing the advantage of standalone versus integrated channels (less interference with instruction processing), since each channel group in a 3033 Processor has its own microcode and arithmetic logic unit and interferes with instruction processing operations in the same situation as do Model 168 channels.

Availability improvements for the 3033 Processor should result from the fact that most fault location programs for a director can be

executed concurrently with normal system operation (the operation of fault locating programs for 2860, 2870, and 2880 channels requires a dedicated processor) and engineering changes can be made more quickly (by reloading director control storage). Serviceability is also improved by the logging of more comprehensive logout data in the I/O extended logout area in the 3033 Processor.

CHANNEL BUS CONTROLLER

Communication between a channel group in a 3033 Processor and processor storage is controlled by the director associated with the channel group and the channel bus controller (CBC), which is contained in the processor storage control function.

The CBC contains the buffers used by the channel groups. One channel request for processor storage can be sent to the CBC at a time. When all directors have a channel request outstanding at the same time, director 1 is given priority over director 2 and director 2 over director 3 for access to the CBC.

Data from processor storage (fetched as a result of a channel write request) and I/O status information is sent to the directors via the CBC. When the CBC contains data for all directors, director 1 has priority over director 2 and director 2 has priority over director 3 for receiving data from the CBC.

I/O ERROR HANDLING

The 3033 Processor has I/O error handling capabilities similar to those provided in other System/370 processors addition to new recovery functions: (1) a machine check external damage code in the fixed logout area to report a channel group error and an associated validity bit in the machine check code in addition to other external damage conditions, (2) an interface inoperable (hung I/O interface) bit in the limited channel logout word, (3) a CLEAR CHANNEL instruction, and (4) an I/O extended logout validity bit in the limited channel logout word.

Channel retry, command retry, and instruction retry facilities, like those provided in other System/370 processors, are provided to reduce the number of abnormal program terminations and unscheduled processor halts that occur because of I/O errors. An I/O extended channel logout is also implemented in the 3033 Processor.

Channel Retry

This feature has been implemented to ensure that most failing channel operations can be retried by error handling routines. When a channel error or a processor error associated with a channel operation occurs, the channel status word (CSW) and a limited channel logout word are stored in the fixed lower storage area (I/O communications area) during the I/O interruption or instruction retry. The limited channel logout word contains processor-independent data and provides more exacting status information about the channel failure. The CCH routine passes this data to a device-independent error recovery routine to be used in the retry of the failing 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 interruption and programmed

error recovery are not required. An indication is presented when the control unit recognizes this situation. Only block multiplexer channels can perform a command retry. The command retry feature is implemented in the control units of 3330-series, 3350, and 2305 disk storage.

Instruction 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 execution function determines whether or not the instruction retry threshold for the particular instruction has been passed. If instruction execution has not gone beyond predetermined points, the instruction is retried automatically by the instruction retry function without programming assistance. A machine check interruption 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 interruption 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, an instruction retry is possible.

Additional Recovery Features

An I/O extended channel logout of 576 bytes is implemented in the 3033 Processor. When a channel control or interface control check occurs for a byte or block multiplexer channel, the CSW, limited channel logout, and extended channel logout are stored during the I/O interruption or instruction retry. Bit 15 in the limited channel logout indicates whether the extended channel logout is valid. The extended channel logout is stored beginning at the location specified in word 172 in the fixed area if the I/O extended log mask (bit 2 in control register 14) is one.

Two additional channel error conditions and an additional I/O instruction are implemented for the byte and block multiplexer channels in the 3033 Processor. These are the interface inoperative (hung I/O interface) condition, channel-not-operational (hard channel group error) condition, and privileged CLEAR CHANNEL (CLRCH) instruction.

An interface inoperative condition indicates a channel error. It exists when a malfunction of an I/O interface persists after selective reset is signaled on the interface. This error condition causes an I/O interruption and the storing of the CSW, limited channel logout, and I/O extended channel logout. Interface control check is indicated in the CSW and interface inoperative is indicated in bit 27 in the limited channel logout. The logout valid bit is one in the limited channel logout if the I/O extended logout is valid.

A channel-not-operational condition indicates a channel group error. It exists when one or more channels in a group have entered the not-operational state without performing an I/O system reset on the I/O interface. That is, the error on one or more channels is severe enough to preclude continued operation of the channels in the group. This error condition causes an external damage machine check interruption and storing of the external damage code followed by an I/O interruption and the storing of the CSW, limited channel logout, and I/O extended channel logout. The external damage code indicates a channel-not-operational condition in bit 3. The CSW indicates a channel control check and the logout valid bit (15) is on in the limited channel logout when a valid I/O extended logout was also stored.

Recovery procedures for both these channel error conditions can be implemented using the CLRCH instruction. The CLRCH instruction causes an I/O system reset to be attempted for the addressed channel. The condition code indicates whether or not the I/O system reset was performed (0 indicates a successful reset, 3 indicates the channel does not exist or is offline). A timeout interruption occurs if the CLRCH instruction is not executed within two seconds.

When an interface inoperative condition exists on a channel, the CLEAR I/O instruction should be issued to all active subchannels before the CLRCH instruction is issued to avoid erroneous continued operation of the channel after the CLRCH instruction is issued. When the condition code presented for an CLRCH instruction indicates an I/O system reset was completed on the channel, I/O operations that were in progress on the channel can be restarted and system operation can continue without a stop and re-IPL.

When a channel-not-operational condition exists for a channel group, the CLRCH instruction can be issued to each of the channels in the group. This causes an I/O system reset on each channel followed by an initial microprogram load of the microprogram for the channel group. This re-IMPL (which is performed automatically only after the CLRCH instruction has been issued to each channel in a group after the time interval required for a re-IMPL elapses) can cause the channel-not-operational condition for the channel group to be reset. If it is (as indicated by issuing a TCH to each channel in the group after the time interval required for a re-IMPL elapses), I/O operations in progress on the reset channel group can be restarted and a processor termination and re-IPL are avoided.

The extended channel logout consists of detailed status information about the director and all six channels in the affected group at the time the channel error occurred. This data should be logged (in SYS1.LOGREC) for processing by the Processor Logout Analysis program.

10:25 MAINTENANCE AND RETRY FUNCTION

The maintenance and retry function provides the data path between the 3036 Console and other components of the 3033 Processor for manual and service operations performed using the console. This function uses the control storage and circuitry of the execution function, an instruction counter, additional registers, a maintenance control word register, and special data paths to the 3036 Console.

Certain of the additional registers are used to handle data that passes between the console and processor components during manual and service operations. Certain other registers are used to buffer information required by the execution function to perform instruction retry and continue processing following a successful retry. The maintenance control word (MCW) register, together with an MCW latch and MCW decremter, are used to implement the functions of the DIAGNOSE instruction. In addition, the trace unit receives status information from the maintenance and retry function.

TRACE UNIT

The trace unit receives certain status data (digital information) from the instruction processor and channels during operation of the 3033 Processor, stores the data in trace buffers, and when a predetermined event occurs presents the trace buffer data to a console processor in the 3036 Console (normally that of the service support console) for storing on its diskette drive.

The trace unit obtains the following data:

- Information from 191 fixed lines (permanently monitored points) in the 3033 Processor. Every processor cycle (57 nanoseconds), the data from these 191 lines is placed in a trace buffer. The buffer has a maximum capacity of 32 processor cycles of data. A wraparound technique is used to store data in the trace buffer so that the buffer always contains information regarding the last 32 processor cycles.
- Information from eight movable probes in the 3033 Processor that the customer engineer can place where desired. Information from these probes is placed in a second trace buffer eight times every processor cycle. The maximum capacity of this buffer is 32 processor cycles, each with all probes recorded eight times per cycle. A wraparound technique is also used to retain data from the most recent 32 processor cycles in this buffer.

Information from the 191 fixed points and eight probes is stored continuously in the trace buffers in wraparound fashion until a predetermined event occurs. When the event is recognized, recording stops until a console processor (normally that of the service support console) polls the trace unit to determine whether it has new data to be written to the diskette drive. If so, the trace data is then sent to the console processor, which formats the data, time stamps it, and writes it to its diskette drive. Trace unit recording resumes as soon as the trace data has been transferred to the console processor.

The event that is to cause the trace data to be transferred to the console processor is established using the console (mode control frame) and can be one of the following:

- Machine check error detected in the execution function, IPPF, PSCF, or processor storage
- Control storage address compare
- Processor storage address compare
- Control storage and processor storage address compare
- Hang detect condition
- Remote signal (one of the above five conditions, as selected, occurs in the other processor in a tightly coupled multiprocessing configuration)
- A condition established by the customer engineer or a logic level transition to minus or plus (as seen through high-speed probe 1) or a machine check error detected in the execution function, IPPF, PSCF, or processor storage

When the predetermined event occurs, a time interval is established. If a machine check (processor) logout, I/O extended channel logout, or power logout occurs during the interval, it is assumed to be associated with the event that caused this recording to take place. The logout data is then formatted as appropriate to its type, time stamped, and written to the diskette drive. The logout is associated with the trace event and can be displayed when the trace data is displayed. If a logout does not occur within the established interval, no logout is associated with the recorded trace event.

The processor logout data is divided into three areas for the purpose of recording: status area (corresponding to the fixed logout area from processor storage locations 0 to 191), local store area (corresponding

to the fixed logout area between locations 216 and 511), and the processor area (corresponding to the model-dependent machine check extended logout area beginning at the location indicated in control register 15).

A maximum of twelve trace events and their associated logouts can be recorded on the diskette on a diskette drive. The mode control frame is also used to establish the trace recording mode, which determines what is done when twelve trace events have been recorded. Recording can continue, with each new trace event replacing the oldest recorded trace event (wrap mode) or recording can stop after twelve events are stored (freeze mode).

A list of the trace events recorded on the diskette drive can be displayed on a CRT of the 3036 Console using the event frame. The trace display frame can then be utilized to display a specific trace event record and its logout data. Trace unit data will be of particular value in diagnosing malfunctions in the IPPF and execution function, since the operation of these functions is overlapped.

SECTION 20: THE 3036 CONSOLE FOR THE 3033 PROCESSOR COMPLEX

20:05 COMPONENTS, FUNCTIONS, AND GENERAL OPERATION

The standalone 3036 Console provides operational and maintenance functions for the 3033 Processor. The 3036 Console embodies a new design that enables maintenance operations to be performed concurrently with production processing using two physically and functionally separate operating stations. The dual station design also offers new serviceability facilities and console backup. If one station is not functional, the other can be used to continue normal processing.

The technology used in the 3036 Console is MST, instead of the SLT used in the 3066 Console for the Model 168. In addition, the cables connecting the 3036 Console and 3033 Processor can be up to 70 feet in length, instead of up to 30 feet, as for the 3066 Console and 3168 Processing Unit. Note that the 2150 Console can be attached to the 3033 Processor Complex as an additional console. However, the Remote Operator Console Panel feature (which permits duplication of the operator controls of the standard console on the 2150 for some System/370 processors) is not supported when the 2150 is attached to a 3033 Processor.

RPQs are available that permit a 3056 Remote System Console to be attached to the 3036 Console and located up to 150 feet from the 3036. The 3056 parallels the functions of console station A or B under microcode control. All display frames and keyboard operations available at the enabled console station are also available to the 3056. Control panel functions, such as console IMPL, set time-of-day clock and console power on/off, are not available at the 3056.

COMPONENTS

As shown in Figure 20.05.1, the 3036 Console is an L-shaped unit that has a control panel and two operating stations, each of which includes a cathode ray tube (CRT) display and keyboard. The rightmost station, as shown in the figure, is designated as console station A while the other is designated as console station B. The CRT and keyboard of both console stations can be rotated to the right 190 degrees.

The rotational capability of both displays allows either console to be positioned for operation of both stations by one operator from one position or separate operation of the stations by two operators from positions that do not interfere with each other. Leg room is provided behind that shown for console station B for use when the CRT and keyboard have been rotated to the right.

The activity to be performed by each operating station, operational or service functions, is initially selected at IMPL time using the control panel on the 3036. That station to be used primarily to communicate with the operating system to control normal operations is designated as the operator station, while the other station is designated as the service support station and is used to perform service operations.

Either operating station can be designated as the operator or service support station. However, normally station A is the operator station and station B is the service support station. The initial console function designations can be altered later if necessary. The service support station can also be defined as an alternate or additional console to the operating system being used. If the processor must be dedicated to maintenance functions, both operating stations can be used for service operations.



Figure 20.05.1. The 3036 Console

Each operating station of the 3036 Console contains the following components:

- A microprogram-controlled console processor that controls the operation of the station
- A dedicated CRT display
- A dedicated keyboard
- A dedicated standard I/O interface that is used to connect a component of the operating station to a channel in the 3033 Processor
- A diskette drive that reads and writes removable diskettes. Microcode and microdiagnostics for the 3033 Processor and 3036 Console are provided on diskettes. The diskette containing the microcode for the 3033 and 3036 is called the system diskette. A diskette drive can be configured for access by either console processor. Another mode of operation of a diskette drive enables certain data contained on a mounted diskette to be sent to the 3033 Processor.

In addition to two operating stations and the control panel, the 3036 Console contains a port selector that is used to directly connect the two console processors to their diskette drives and other components of the processor, one 1200-baud modem that can be used to connect console station B to a communications line for remote maintenance operations, and one counter that is used by both console processors for time stamping operations. If the remote maintenance capabilities of the 3036 Console are to be utilized, a communications line must be installed for this purpose.

Each of the two console stations of a 3036 Console attaches to a channel in the 3033 Processor. A console station can be attached to a byte or block multiplexer channel but normally would be attached to a byte multiplexer channel. Each console station uses one control unit position on the channel to which it is attached and the two console stations can be attached to different channels and channel groups. For availability purposes, the two consoles should be attached to different channel groups.

Console station A has two different I/O device addresses while console station B can have three different I/O device addresses. The device addresses can be any of the possible 256 available for a channel. One address for each console station is assigned to the CRT and keyboard for use with the program frame. The second address for each console station is assigned to the diskette drive for use when it is accessed by a program executing in the 3033 Processor. The optional third address for console station B is used to address an emulated 2955 Remote Analysis Unit during a remote maintenance operation.

Only one of the I/O units for each console station can be accessed by a channel at a time via the associated standard I/O interface for the station. Thus, for example, when the CRT and keyboard of a station are being used for communication with the operating system, the diskette drive of that station cannot be accessed by a program in the 3033 Processor.

FUNCTIONS AND GENERAL OPERATION

Functions

The 3036 Console is used to perform the following major functions:

- Sequence, monitor, and control power
- Load instruction processor (execution function), director, and console microcode
- Configure certain processor components and 3036 Console functions
- Control 3033 Processor operations (IPL, start, stop, etc.) and communicate with the operating system being used
- Configure the 3036 Console for concurrent maintenance operations
- Exercise processor components to locate a malfunction
- Display processor and channel indicators and logouts
- Execute microdiagnostics for the 3036 and 3033 that are contained on diskettes
- Perform remote maintenance operations

The functions listed are performed using the two CRT displays and keyboards and the control panel. The CRTs and keyboards operate in display mode only (there is no printer-keyboard mode, as is supported for the display consoles of certain other System/370 processors). A specific hard-copy printer is not available for the 3036 Console. Hard copy can be obtained using the hard-copy support provided by the operating systems that support the 3033 Processor (OS/VS DDOCS, for example).

A system status recording (SSR) facility is also implemented in the 3036 Console. This microcoded function monitors the occurrence of unusual operational events that are not recorded by recovery management routines. The events that are monitored are console power on, console power off, meter key switch on, disabled wait state, check stop, processor IMPL, IPL, start, restart, stop, alter, system reset, CPU reset, hang detect condition, SSR enabled, and SSR disabled. The preceding are recorded when the meter key switch is in the off position. Only the meter key switch off condition is recorded when the meter key switch is on.

When an event occurs, the console processor of the service support console formats an SSR record, time stamps it, and writes it to its diskette drive. Up to 680 system status records can be recorded on an system diskette. After 680 events are recorded, each successive record replaces the oldest record. These records can be displayed on the CRT using the system status recording frame. A maximum of 20 events can be displayed at a time.

The SSR data records provide a history of the unusual operations that were performed by the operator and indicate when the customer engineer had control of the processor. The SSR data and that recorded in SYS1.LOGREC provide a chronological record of all system events.

General Operation

The images that are displayed on the CRTs are called frames. The set of display frames provided is used to control both normal and service operations. Each frame indicates functions that can be selected and/or provides a specific display. Frame selection and the function or display to be performed are accomplished using the keyboard. The code associated with the frame or function must be keyed in. There is no light pen.

Frame selection and display for a CRT are controlled by its associated console processor. The two console processors operate independently from each other; however, they can communicate with each other. Communication between the two console processors in the 3036 Console is accomplished using a general purpose communication register.

Communication is required, for example, whenever configuration data is entered via a console station, since this data must be written to the system diskette on both diskette drives. Communication is also required to maintain the status information that appears on the last line of certain frames whenever the frame being displayed on each CRT contains the status line.

The console processor uses the counter in the 3036 Console for time stamping events, such as processor logouts, channel logouts, trace events, power logouts, and SSR records. This counter has a one-second resolution and is powered by a battery when power in the 3036 Console is off. The counter can operate for approximately 34 years before overflowing.

20:10 CONTROL PANEL

The control panel located on the 3036 Console, shown in Figure 20.10.1, provides basic power control for the 3033 Processor Complex, controls for configuring the two operating stations, console microcode IMPL controls, diagnostic on IMPL switches, I/O interface enable/disable switches, usage meters, the time-of-day clock enable set switch, the means to activate telecommunications operations, audible alarm volume control, and error indicator panels. Other operational control switches are contained on the keyboards.




Figure 20.10.1. Control panel on the 3036 Console

The control panel contains the following:

- Power on and power off pushbuttons, power select switch, and power off pending indicator. When the power on pushbutton is pressed after first pressing the IPC reset pushbutton, the initial power controller (IPC) in the 3036 Console turns on power in the 3036 Console. Microcode for both console processors is loaded after the 3036 Console is powered on. The IPC monitors power after the 3036 Console has assumed power control (power monitoring microcode is loaded and active) to ensure that the console microcode is reacting to power faults in time.

No more powering occurs if the power select switch is set to the console position. If this switch is set to the system position, a power on of the components in the 3037 Power and Coolant Distribution Unit, 3033 Processor, and I/O devices set to remote power occurs after the power on and IMPL of the 3036 Console are completed. When the processor power-on sequence is completed successfully, an IMPL of execution function and director microcode is automatically performed in the 3033.

The power-on sequence is microprogram-controlled by the console processor of the operator station. Components are powered on in the following sequence: coolant distribution unit, all 3033 Processor components except the directors and channel groups, director/channel group 1, director/channel group 2, director/channel 3, if present, and all I/O devices that have their power control switches in the remote position. Note that the 3037 and each director/channel group has a local/remote power control switch and is powered on during the power-on sequence only when the switch is in the remote position. Thus, these units can also be powered on and off using their local power on and off pushbuttons.

Components can also be powered on and off via the 3036 Console using the power control frame. The following units can be separately powered on and off: the coolant distribution unit, processor storage, instruction processor function and high-speed buffer, each channel group, and I/O devices. Separate powering of these units enables card replacement to be performed without powering down the entire processor complex.

If any component does not power on properly, the power-on sequence stops at that point and the power control frame is displayed and indicates an error condition. At the completion of a successful power-on sequence, the configuration frame is displayed on the service support station and the program frame is displayed on the operator station.

The power on pushbutton backlight is red while the power-on sequence is performed. When all power supplies, as specified by the power select switch, are up and no check condition exists, the power on pushbutton backlight turns white.

When the power off pushbutton is pressed, the power-off sequence is initiated for the 3033 Processor, 3037 Power and Coolant Distribution Unit, 3036 Console, and I/O devices with their power control switch in the remote position regardless of the setting of the power select switch. Since processor storage and all control storages are implemented in monolithic technology, the contents of processor storage, as well as its storage protect keys and control storages, are lost when a power off is performed. The power off pending indicator lights when the power off pushbutton is pressed.

- IMPL pending and microcode power control indicators. After a 3036 Console power-on sequence is completed, the IMPL pending indicator

is lit while an IMPL is being performed for both console processors. An IMPL frame is displayed on the CRT while an IMPL is being performed for a console processor. This frame indicates the IMPL steps that have completed, the step currently executing, and the steps that are not applicable to this type of IMPL (power-on IMPL or pushbutton IMPL).

When the power select switch is set to the system position, the microcode power control indicator is lit after the power-on sequence is performed in the 3033 Processor. The microcode power control light indicates the 3036 Console has control of power and remains on until a power off occurs.

- IPC reset pushbutton. This pushbutton must be pressed before a power on can be performed. It resets any console power fault conditions. Console power fault conditions are indicated in lights on a console IPC power fault board located on a swinging panel at the left end of the console panel. The IPC reset pushbutton turns off these fault indicator lights.
- Unit emergency switch. When this switch is moved to the down (power off) position, all power is removed from the 3036 Console and 3033 Processor immediately (that is, without sequencing). Power remains on in the power distribution unit. When in the down position, the unit emergency switch can be returned to the up (power enable) position only by service personnel.
- I/O interface switches (one for console station A and one for console station B). When an I/O interface switch is in the enable position, the connection between its associated console and a channel in the 3033 Processor via the standard I/O interface is enabled. This connection is disabled when the switch is in the disable position. The I/O interface switch must be in the enable position when its associated CRT display and keyboard are to be used for operator-to-operating system communication (via the program frame), the diskette drive is to be accessed via a program in the 3033 Processor, or emulation of the 2955 is to occur.
- Console IMPL (IMPL A and IMPL B) pushbuttons. When a console IMPL pushbutton is pressed, an IMPL is performed in the corresponding console station if the security key on the associated CRT is in the horizontal position. IMPL is not performed if the security key is in the vertical position. These pushbuttons do not also cause processor and director microcode to be loaded in the 3033. The security switches on the two CRT displays need not be turned on to perform a power-on IMPL. A pushbutton-initiated IMPL of one station should not be initiated while the other station is executing an IMPL. When the system diskette is mounted for the console station being IMPLed, the basic console diagnostics on the diskette are executed before the IMPL is initiated.
- Diagnostic on IMPL switches (one for each console station). This switch setting is inspected only if the console and power diagnostic diskette is mounted on the drive for the console station being IMPLed. When the console and power diagnostic diskette is mounted, this switch determines whether the console or the power diagnostics are loaded and executed. When the switch is in the enable position, console diagnostics are executed.
- Operator console on IMPL switch. This is a two-position switch that is used at three different times. It is used during a power-on IMPL to assign the function to be performed (operator or service support) to the two operating stations. Station A can be assigned as the operator station and station B can be assigned as the service

support station (this is the normal configuration) or this assignment can be reversed.

When a pushbutton IMPL is performed and the switch is set to the console station being IMPLIED, a local IMPL of that station occurs as usual. Otherwise the operator is prompted to indicate whether full console reconfiguration is to be performed.

When the operator uses the configuration frame to specify the mode in which the 3036 Console is to operate (normal or maintenance), the switch setting determines which console station is to be assigned the operator function.

- TP active/key reset pushbutton and activate TP keyhole. In order to select remote service functions using the TP link frame, telecommunications (the modem in the 3036 Console) must be activated. Activation is accomplished by inserting the CE key in the activate TP keyhole and turning it. The TP active/key reset pushbutton is yellow during the time telecommunications is active and the CE key can be removed any time thereafter. Telecommunications can be deactivated at any time by pressing the TP active/key reset pushbutton.
- Time-of-day clock enable toggle switch. This switch must be held in the enable set position in order for the time-of-day clock to be set using a SET CLOCK instruction.
- Alarm volume control switch. This switch can be rotated to control the level of the program-controlled audible alarm. This switch controls the audible alarm in the 3036 Console that is sounded by the issuing of a sound alarm command to the program frame unit address.
- Customer meter, CE meter, and meter keyhole. The meter to be used is selected by inserting a CE key in the meter keyhole and setting the key to the customer or CE meter position.
- Console Processor Error Indicator Panel. This panel is on a movable box stored inside the control panel behind the door on the left-hand side of the panel. The top half indicates errors associated with console station A while the bottom indicates errors associated with console station B.
- Initial Power Control (IPC) Fault Indicator Panel. This panel is located on the door that covers the console processor error indicator panel. This panel is provided to aid in isolating power faults that may occur in the 3036 Console during the power-on sequence.

20:15 CRT DISPLAY AND KEYBOARD

The CRT display for each operating station can display 25 lines of 80 characters each simultaneously. When the program frame is active, physically the CRT and keyboard appear as a 3277 Display Station Model 2 attached to a 3272 Control Unit. A subset of the commands and orders for a 3277 operating in display mode are used to control the CRT.

Two new commands are provided for the CRTs on a 3036 Console. One new command sounds a loud alarm, which is different from the audible alarm that is sounded when a specific bit is on in the write control character of a write command for a 3277 display. The alarm command is not immediate and, therefore, should be issued with the suppress incorrect length (SILI) bit on in the UCW. The word ALARM is also written on line 25 of the display. The other new command is a sense I/O command. The device-dependent error recovery procedures used for the 3277 Model 2 can be used for the CRTs of the 3036 Console.

In addition to intensity control (provided to vary the intensity of displayed characters) and bezel indicators (that indicate a console or processor error condition), each CRT has a security key, which is located on the right side of the unit. This key is used to determine whether or not IMPL, frame selection, or an external interruption can be performed using the keyboard. When the security key is in the vertical position, it can be removed. In the horizontal position, the key cannot be removed. (See discussion in Section 20:25 for the frames available to each type console.)

Operators and service personnel perform functions not provided by the control panel on the 3036 console using the keyboard provided for each operating station. The functions available are grouped into display frames that have two-character identifications. The operator or service person selects the desired frame by keying in the appropriate frame identification, after pressing the SEL FRAME key, and selects the desired function by keying in the identification assigned to that function, as shown on the selected frame. The associated console processor performs the actions required to display the requested frame and execute the requested function.

The keyboard provided for each operating station contains 26 alphabetic character keys (both upper and lower case), 10 numeric keys, 26 special character keys, 8 keys for cursor control, and 8 keys for special functions. In addition to the shift, lock, cancel, keyboard reset, enter, and spacebar function keys (which essentially control keyboard operations), the following function keys are provided:

- START and STOP to start processor operations when the processor is in the manual (stopped) state and stop processor operations (place the processor in the manual state).
- IRPT to present an external interruption to the processor. When pressed, the key causes the position of the security key on the side of the display to be tested. If the key is in the vertical (secure) position, an interruption is not presented and the alarm is sounded. Otherwise, the interruption is presented.
- PFK SEL to change keys 1 through 9, and 0, and the dash and ampersand keys to program function keys 1 through 12, respectively. The letters PFK are displayed at the bottom of the active frame (valid only when the program frame is active).
- SEL FRAME to display a desired frame. After this key is pressed and the identification of the desired frame is keyed in, the selected frame is displayed. Note that the frames that can be selected depend on whether the operating station is an operator or service support station and the setting of the security key on the CRT unit, as previously discussed.

While these functions are typical, they do not necessarily apply to all frames.

20:20 DISKETTE DRIVE AND DISKETTES

A read/write diskette drive is used to load microcode and microdiagnostics and to record trace events, SSR records, and logout data. The diskette drive in the 3036 Console is different from the diskette drive used for most other System/370 processors (3115 to 3168 processors). Both sides of the diskette used on the diskette drive in the 3036 Console are utilized and the diskette has a larger capacity than other diskettes. All even-numbered tracks are contained on one side while all odd-numbered tracks are contained on the other.

All the microcode required by the 3033 Processor (execution function and directors) and 3036 Console can be contained on one diskette. Three such system diskettes are sent to each 3033 Processor installation. The same system diskettes are sent to all 3033 installations. Thus, a system diskette used in one 3033 installation can be used at another 3033 installation as long as the diskette and the other 3033 Processor are at the same engineering change level.

The first time a given system diskette is used for a remote maintenance function, the desired 3033 Processor serial number, along with other information, should be written on the diskette, using the TP link frame. The processor serial number (obtained using a STORE CPU ID instruction) is not checked against the serial number recorded on the mounted system diskette during the IMPL operation.

In addition to microcode, a system diskette contains formatting and spacing information for editing MCH and CCH logout data (called frame verbiage), basic console microdiagnostics, and areas for logout data. Processor logout data (that is logged after a machine check), trace unit data, director logout (I/O extended logout) data, and system status recording unit data, are written on the system diskette mounted on the diskette drive (normally the diskette drive for the service support station). Power logout data is normally written to the diskette drive of the operator station.

A verbiage frame is a representation of a screen image consisting of text and scan buffer codes (format and spacing information) for an MCH or CCH logout record type. This frame data is utilized by the EREP program. Frame information for the logout data of the 3033 Processor is structured to have a definition that is consistent with that for the logout data of the 3032 and 3031 Processors. This frame approach makes the EREP program processor-independent among 3033, 3032, and 3031 Processors and insensitive to ECs, since only the frame information need be changed when an EC occurs.

A system diskette can contain up to twelve processor logouts, up to twelve director logouts for each director, up to twelve trace unit records, up to 680 SSR records, and one full track of power logout records at a time. A system diskette also contains a time-of-day value and a configuration record that reflects the assignments made using the configuration frame.

The first time a system diskette is used after the IMPL procedure, a frame appears to prompt the customer engineer to enter the date and time of day. When this information is entered, the console processor writes the value of the counter, date, and time of day on the operational diskette. The date and time are also saved in console processor control storage.

When time stamping is required, the console processor adds the difference between the current counter value and the saved counter value to the saved time to obtain the current time. Whenever an IMPL is performed, the date and time recorded on the operational diskette are read into console processor control storage for use in time stamping.

Power logouts are automatically written to the system diskette. Processor logouts, director logouts, trace events, and SSR records are written to the system diskette only if recording mode is enabled for these logouts using the mode control frame. The recording of logouts from each director is individually controlled and the writing of all logouts (processor and logouts from all three directors) can be disabled.

The advantages of writing processor and director logouts to the system diskette, as well as to SYS1.LOGREC, is that the customer

engineer can access these records utilizing the service support station without interfering with normal operations and without accessing the SYS1.LOGREC data set to obtain logouts associated with trace events. The customer engineer can display logout data or it can be sent to a remote location for analysis concurrently with normal operations.

During normal system operation, a system diskette must be present on the diskette drive for both operating stations. The third system diskette is provided for backup. Other diskettes containing diagnostics are also sent to each 3033 installation. Three processor diagnostic diskettes, one director diagnostic diskette, and two console and power diagnostic diskettes are provided. Console, processor, power, and director diagnostics are executed under the control of the CRT and keyboard using service frames.

Normally a diskette drive can be accessed only by its associated console processor and not by a program executing in the 3033 Processor. However, the configuration frame can be used to place a diskette drive in service record file mode. This mode establishes a connection between the diskette drive and a channel via the standard I/O interface associated with the diskette drive. When this mode is enabled, certain information contained on the mounted diskette, such as logout data, can be obtained by a program executing in the 3033 Processor. Service record file mode is designed to be used by customer engineers to access the logout and status data via an OLT, for example.

Note that a program interfaces with a diskette drive via the console processor. That is, a program cannot issue START I/O instructions directly to the diskette drive to read and write data. All diskette reading and writing operations are performed by the console processor and there is no way a user-written program executing in the 3033 Processor can cause data to be written to a diskette drive.

20:25 DISPLAY FRAMES

The display frames provided for the 3036 Console are grouped according to the basic functions they provide. There are operator and service frames.

A list of the frames provided and their two-character identification is contained on the index frame. When the index frame is displayed on the CRT of a console station, the two-character identification of each frame that is not usable by that console station based on its assigned function is not displayed.

For example, when the index frame is displayed on the operator station, the MD identification is not displayed opposite the MICRODIAGNOSTIC frame listing, since this frame is not available to the operator station.

The operator frames are index, alter/display, configuration, director configuration, operator, program, system activity, and system status recording. Service frames are clock margin, CPU CS alter/display, director service, event index, log index, mode control, microdiagnostics, power control, processor service, trace display, TP link, voltage margin, director indicator/logout, instruction preprocessing function indicator/logout, execution function indicator/logout, maintenance indicator/logout, processor storage control function indicator/logout, log, and microdiagnostic control frames.

As indicated previously, frame selection is accomplished using the SEL FRAME key and the two-character identification of the desired frame,

under the control of the security key on the CRT. Any frame can be selected regardless of the frame currently being displayed.

When the security key for the operator station is in the vertical (on) position (or removed), no frame selection is permitted and the frame being displayed when the key was turned is retained. The audible alarm controlled by the write command is sounded if an attempt is made to select another frame. When the security key is in the horizontal (off) position, only the program, index, configuration, power control, or system activity frame can be selected.

When the security key for the service support station is in the vertical (on) position (or removed), no frame changes can be made. The frame being displayed when the key was turned is retained. When the security key is in the horizontal (off) position, frame selection can be made and any frame except program (when service record file or remote program mode is active for the console) or power control can be selected from the service support station.

OPERATOR FRAMES

Operator Frame

The operator frame is shown in Figure 20.25.1. It is used to control basic operations of the processor, such as IPLs, resets, restarts, address compare stops, store status function, etc. The store status function provides the same information for the 3033 Processor as for other System/370 processors.

The operator frame is also used to IMPL and initialize certain processor components. The following can be IMPLed and initialized: execution function and all directors present, execution function only, or the specified director. Initialization consists of establishing communication between the unit(s) being IMPLed and the console. An IMPL of execution function microcode only without initialization can also be performed.

The action the processor is to take when a machine check occurs, stop or continue processing, is also set using the operator frame. When set to stop, the processor enters the stopped state immediately after a machine check occurs. There is no logout to processor storage or the diskette. After such a stop, the customer engineer can cause a logout to the diskette using manual procedures.

The current PSW, system indicators, and console station status are displayed at the bottom of the operator frame. This line of status is also displayed on all frames that can be selected when the system diskette is mounted. When the 3036 Console is in normal mode, line 25 information is displayed on both CRTs. When maintenance mode is in effect, line 25 information is displayed only the CRT of the operator station.

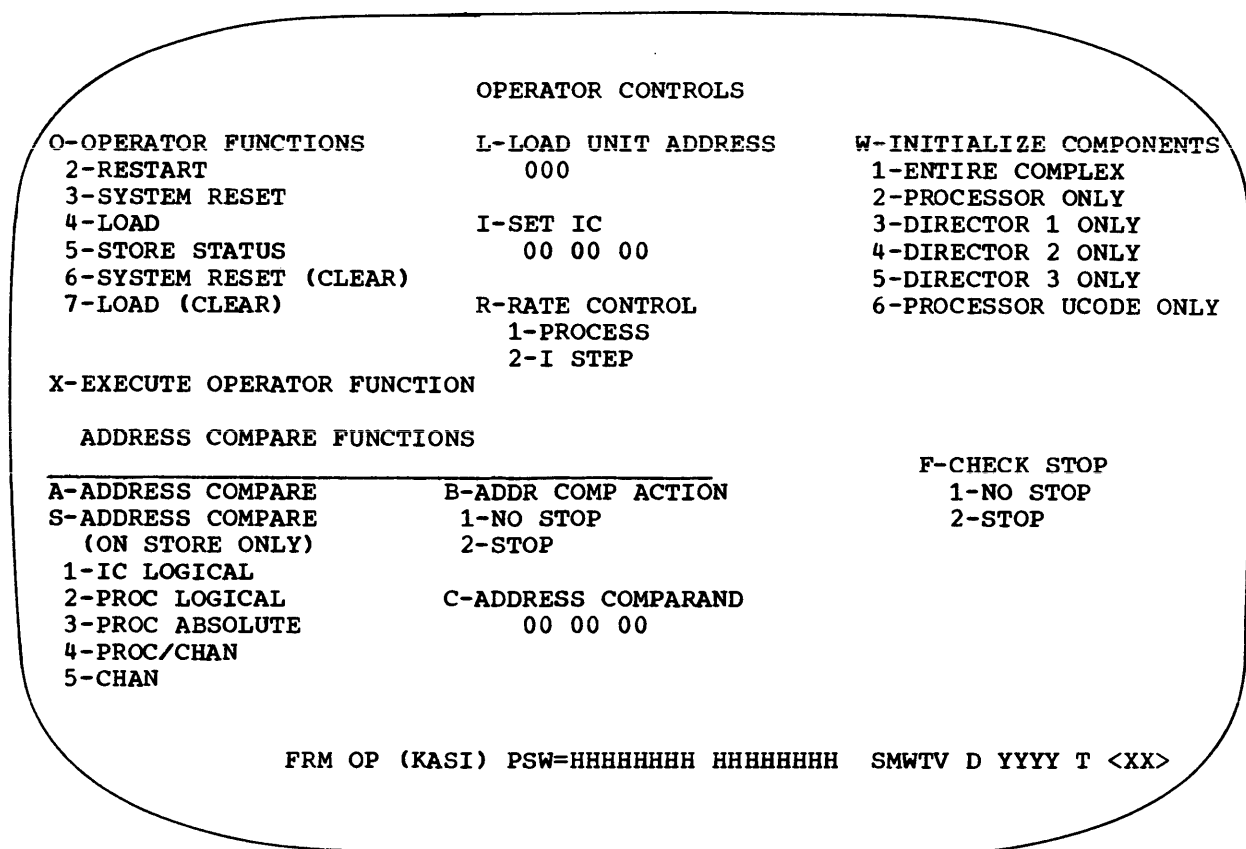


Figure 20.25.1. The operator frame

Program Frame

The program frame is shown in Figure 20.25.2. This frame is used for communication between the operator and the operating system during processor operation. The layout shown is that supported by DIDOCS in OS/VS. The program frame is automatically displayed on the operator console at the successful completion of an IPL and also can be selected using the SEL FRAME key. The program frame is cleared and initialized (all messages erased) when any one of the following is performed using the operator frame: system reset, system reset and clear, load, and load and clear. The program frame is also cleared and initialized by an IMPL, an error IMPL, an I/O system reset or an I/O selective reset not manually initiated, or activation of the program frame from the C1 configuration frame.

When the program frame is displayed, the PFK SEL key can be used to establish keys 1 to 0, -, and & as twelve program function keys. When any one of these keys or the KEYBD RESET key is pressed, these keys return to their normal function (that is, program function mode is no longer active).

Alter/Display Frame

The alter/display frame is shown in Figure 20.25.3. It enables an operator to alter and/or display the facilities listed on the frame. The active subchannels can be displayed only. The processor must be stopped (in manual mode) in order to alter or display data.

SYSTEM MESSAGE AREA, LINES 1-19

BLANK

INSTRUCTIONS

ENTRY AREA

ENTRY AREA

WARNINGS

SYSTEM AVAILABLE (KASI) PSW=HHHHHHHH HHHHHHHH SMWTV D YYYY T <XX>

Figure 20.25.2. The program frame

Configuration Frame

The configuration (C1) frame is shown in Figure 20.25.4. It should be selected when a configuration change is to be made. This frame automatically appears on the service support station after an IMPL or power-on IMPL is performed using the control panel and after IMPLs that take place after certain console errors occur.

The configuration frame displays the 3033 Processor features installed, the currently available processor storage size and address assignments, information about the diskette mounted for the console station, the processor serial number, and the facility currently controlling the I/O interface (service record file, program frame, or remote function). Processor storage reconfiguration, making program frame or service record file mode active, and setting a value in the counter in the 3036 Console are accomplished using this frame. The configuration frame is also used to configure the console stations. See Section 20:30 for a discussion of console configurations.

FUNCTION

- A - ALTER
- D - DISPLAY

FACILITIES

- M - MAIN STORAGE REAL
- V - MAIN STORAGE VIRTUAL
- K - MAIN STORAGE KEYS
- P - PSW
- G - GENERAL REGS
- F - FLOATING POINT REGS
- C - CONTROL REGS
- I - PROCESSOR WORKING STORAGE
- B - CHNL BUMP STOR REAL
- L - CHNL BUMP STOR LOGICAL
- A - ACTIVE SUBCHANNELS
- U - CHNL UCW LOCAL STOR
- D - DIRECTOR LOCAL STOR
- S - CHNL DATA BUFFER L.S.
- T - CHNL STOR ADAPTER BUFFERS

KEY IN ADDRESS

AM 000000

(KASI) PSW=HHHHHHHH HHHHHHHH SMWTV D YYYY T <XX>

Figure 20.25.3. The alter/display frame

- CONFIGURATION -

C-CONSOLE CONFIGURATION

- 1-ADD
- 2-DELETE
- 3-ASSIGN OTHER
- 4-TAKE FROM OTHER
- 5-ENTER MAINT MODE
- 6-ENTER NORMAL MODE

PORTS THIS SIDE

[01234567]

PORTS OTHER SIDE

[01234567]

PROC FEATURES

M-MAIN STORAGE CONFIGURATION

- 1-BUILD ARRAY
- 2-ENTER ARRAY
- 3-8 WAY INTERLEAVE
- 4-SERIAL

PHYSICAL STORAGE = M BYTES

STORAGE ADDRESSING =

CURRENT ARRAY BUILD ARRAY

A-ACTIVATE

- 1-PR FRAME
- 2-SV REC FIL

I/O USER =

TP USER =

S-SET BASE TIME

BASE = 00.000.00:00

TIME = 00.000.00:00:00

DISKETTE =

EC LEVEL = 0000000

TP LEVEL = 00

PROC SER# = 000000

X-EXECUTE OPERATOR FUNCTION

(KASI) PSW=HHHHHHHH HHHHHHHH SMWTV D YYYY T <XX>

Figure 20.25.4. The configuration frame

Director Configuration Frame

The director configuration frame is used to:

- Enable and disable individual directors (and their associated channel group) and channel-to-channel adapters. When the frame is initially displayed, the state (enabled or disabled) of each installed director and channel-to-channel adapter is shown. A power-on enables all installed channel-to-channel adapters.
- Indicate the directors installed.
- Display the current shared control unit assignments for byte and block multiplexer channels.
- Initially assign, add, alter, or delete shared control units. Any assignment causes the shared UCW assignment record for the channels to be read from both system diskettes and rewritten with the altered information.
- Save shared control unit assignments on an system diskette. This function would be performed when a new system diskette with a higher EC level is received at the installation. This function should be invoked while the old system diskette is mounted. A message then appears requesting mounting of the new diskette. When the new diskette is mounted and the enter key is pressed, the shared UCW assignment record is read from the old diskette and written to the new diskette. A message appears requesting an IMPL of the new diskette.

System Status Recording Frame

The system status recording frame is used to display the contents of the system status records currently contained on the system diskette on a diskette drive (usually that for the service support station).

System Activity Frame

The system activity frame provides a graphic display of processor and channel activity. It is intended to provide a relative indication of activity percentages and can be used to detect conditions such as under-utilized channels. The percentages are not intended to be exact and may vary between systems and from one EC level of system diskette to another.

The system activity frame is shown in Figure 20.25.5. Activity is displayed using horizontal bars. The select function on the right of the display is used to indicate the activity to be displayed.

When the PROC/CHAN function is selected, the following is displayed:

- Bar 1 - Processor activity in the PROC MODE selected in item C
- Bar 2 - Channel activity for the first channel selected in item M (channel mask) from left to right
- Bars 3 to 5 - The second through the fourth selected channels
- Bars 6 to 9 - Processor-to-channel overlap for the same channels displayed above

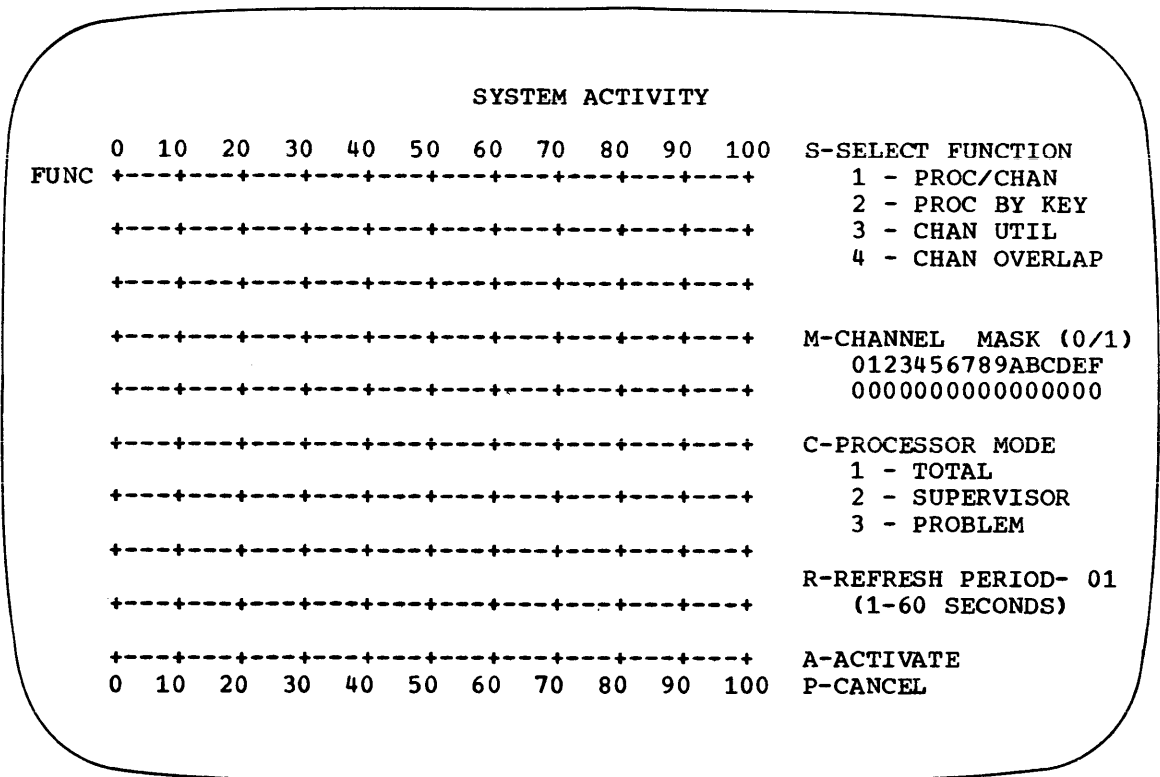


Figure 20.25.5. The system activity frame

When the PROC BY KEY function is selected, the following is displayed:

- Bar 1 - The total activity of the selected storage protect keys (as specified in the PROC KEY MASK (M) for the selected PROC MODE (C).
- Bars 2 to 9 - Processor activity for up to eight selected storage protection keys (as specified left to right in the PROC KEY MASK) for the selected PROC MODE.

When the CHAN UTIL function is selected, the display indicates channel utilization for the first eight selected channels as specified left to right in the CHANNEL MASK (M). The bars show the following:

- Bar 1 - Utilization of any of the selected channels.
- Bars 2 to 9 - Utilization of each of the selected channels.

When the CHAN OVERLAP function is selected, the percentage of overlap for all the channels selected in the CHANNEL MASK (M) item is shown on bar 1.

The refresh period item is used to select the time between display refreshes and from 1 to 60 seconds may be selected. Note that the refresh period entered is a minimum value and the actual refresh period may be longer than indicated.

The ACTIVATE item is used to activate monitoring after all the functions are selected. If the request for the required ports cannot be honored, a message will be displayed in the upper right hand corner of the screen indicating the first unobtainable port (i.e., "PROC PORT BUSY"), otherwise "MEASUREMENT ACTIVE" is displayed. When monitoring is active, a lozenge will be turning on and off to the right of the activate item with every frame refresh.

The CANCEL item is used to stop monitoring. The cancel key also performs this function.

SERVICE FRAMES

Power Control Frame

The power control frame is used to turn power on and off in individual components, as discussed previously in Section 20:10. It is also used to list and clear power logouts and to display power status information.

Mode Control Frame

The mode control frame is provided to enable the customer engineer to control the operation of the trace unit, establish recording or no recording mode for certain logouts, and enable/disable system status recording. The automatic recording of processor logouts, director 1 logouts, director 2 logouts, and director 3 logouts to the diskette drive can be individually enabled and disabled. Whenever the area on the system diskette that is reserved for an individual type of logout becomes full, the next logout of that type overlaps the oldest logout record in the full area so that the area contains the latest twelve logouts. This wrap mode of operation cannot be altered for logout records.

Event Index

The event index frame contains a list of the trace events (a maximum of twelve) currently contained on the system diskette on the diskette drive of the service support station. These are the trace events recorded by the trace unit in the 3033 Processor and associated logout data. The list indicates the type of logout (processor, director, or power) that accompanied the trace event if any, and for a processor logout the logical component (IPPF, PSCF, execution function, or processor storage) in which the error occurred. The customer engineer can then display a specific trace event and its associated logout if any, using the event display frame.

Log Index Frame

The log index frame is used to list the processor and director logouts currently contained on the system diskette on the diskette drive of the service support station. Up to 24 (12 of each type) can be listed at a time. An individual logout can be displayed using an indicator/logout frame.

TP Link Frame

The TP link frame, shown in Figure 20.25.6, is used to initialize the operational diskettes when they are first received in an installation and whenever a new set is sent that contains engineering changes. This frame is also used to perform three different maintenance operations that involve telecommunications and to execute a telecommunications line test. Initialization of the two system diskettes to be used during processing must be done before two of the three remote maintenance operations can be performed.

The initialization function is provided to enter customer identification information (SET ID RECORD function) and security information (SET CUSTOMER SECURITY function) that is recorded on the diskette. The following customer identification is recorded: name and telephone, processor model and serial number, system diskette EC level, and FE branch office and region. The security information that is recorded determines the concurrent maintenance functions that the customer engineer will be allowed to perform.

The maintenance functions that can be performed using the TP link frame are remote program, transmit logs/trace, and remote console. If the transmit logs/trace or remote console maintenance function is not specified on the system diskette, the customer engineer cannot perform that function. After diskette initialization, whenever the TP link frame is displayed an asterisk is displayed to the left of the remote maintenance options that have been authorized by the installation.

```

                                TP FRAME

3033 SERIAL          DISK EC          FE B/O          REGION
CUSTOMER NAME                                PHONE / / EXT

A - INITIALIZE IMPL DISK                      MESSAGE AREA
  1 - SET CUSTOMER SECURITY                    *****
  2 - SET ID RECORD                            *
                                           *
B - REMOTE OPTIONS                             *
  1 - REMOTE PROGRAM                          *
  2 - TRANSMIT LOGS/TRACE                     *
  3 - REMOTE CONSOLE                          *
  (*1)SERVICE RECORD FILE ACTIVE             *
  (*2)PROGRAM FRAME ACTIVE                   *
*1 - MUTUALLY EXCLUSIVE WITH B1,2 and 3      *
*2 - MUTUALLY EXCLUSIVE WITH B1              *
                                           *
C - TP TEST                                    *
  1 - RAISE CARRIER                          *

(KASI) PSW=HHHHHHHH HHHHHHHH SMWTV D YYYY T <XX>

```

Figure 20.25.6. The TP link frame

Before one of the remote maintenance functions can be selected, the CE key must be inserted in the activate TP keyhole on the console panel of the 3036 Console and turned on to activate telecommunications. The CE key can then be removed. Telecommunications can be deactivated at any time by pressing the TP active pushbutton. The remote program, transmit logs/trace, and remote console functions are discussed in Section 20:30.

The TP link frame also indicates whether service record file mode or the program frame is active for the service support station. During a console IMPL, service record file mode is made active for the service

support station by default. That is, microcode to support service record file mode instead of the program frame is loaded.

When the TP link frame indicates the program frame is active, the program frame can be utilized on the service support station and it can be used to communicate with the operating system as an additional console, for example (see discussion of maintenance mode console configurations in Section 20:30).

Microdiagnostics Frame

The microdiagnostics frame is used to initiate and control the execution of the console diagnostics that are contained on the console and power diagnostics diskette. The console diagnostics check console logic, operation of the CRT and keyboard, operation of the modem, and operation of the diskette drive. For certain errors, the suspected failing cards are displayed.

Other Frames

The PROC CS ALT/DISP frame is used to load engineering changes for execution function microcode on the system diskette and display this microcode. Two processor service frames are provided to enable the customer engineer to perform many testing operations on processor components other than the directors, for which the director service frame must be used. Director microcode can be loaded using the director service frame. The clock margin frame allows the customer engineer to vary clock adjustments forward or ahead (cycle duration remains the same) and the voltage margin frame is used to adjust voltages, if required.

The processor diagnostics contained on three of the distributed diskettes are executed using the processor microdiagnostic supervisor frame, and the director microdiagnostic supervisor frame is utilized to execute director diagnostics contained on the other distributed diskette.

INDICATOR/LOGOUT FRAMES

The indicator/logout frames are provided to display logout information from the system diskette or status information about a processor component (director, IPPF, execution function, etc.). The log frame is used to display selected areas of processor logout data (status area, general registers, floating point registers, and control registers).

20:30 CONSOLE CONFIGURATION

TYPES OF CONFIGURATIONS

A console station of a 3036 Console can be configured to perform operator or service support functions and the 3036 Console can be operating in normal or maintenance mode. The state of a console station is defined by combining the 3036 Console mode with the function assigned to the station. The console station states are:

- <NO> - normal mode, operator function. This console station is the operator station. It is used to communicate with the operating system and monitor power.

- <NS> - normal mode, service function. This console station is the service support station. It is used to perform manual operations for event recording and, if necessary, to perform certain maintenance functions.
- <MO> - maintenance mode, operator function. This console station is a combined operator and service support station. It is used to perform the functions described for the <NO> and <NS> states for a console station.
- <MS> - maintenance mode, service function. The console station is used to perform certain service functions. It is considered to be offline to the 3036 Console.

The state of a console station is displayed on the right end of line 25 of its display. When the 3036 Console is in normal mode, one station must be in the <NO> state while the other is in the <NS> state. When the 3036 Console is in maintenance mode, one station must be in <MO> state while the other is in <MS> state. No other state combinations are valid.

The console configuration to be in effect is established initially during IMPL. It can be changed by a re-IMPL or by using the configuration frame. The configuration frame can be used to change the console configuration only if invoked by a console station that is in the <NS> or <MO> state.

The console configuration in effect determines the functions that can be performed by each operating station as well as the 3036 Console and 3033 Processor components that can be directly accessed by each of the two console processors.

A console processor electrically connects to a component in the 3033 Processor or a diskette drive via a common interface called a port. When a console processor and component are port connected, data can be transferred between them. The standard I/O interface to a 3033 Processor channel provides the other means by which certain operating station components (diskette drive, CRT display and keyboard, or 2955 emulator) can be connected to the 3033 Processor for data exchange.

Eight ports, addressed 0 through 7, are defined for the 3036 and 3033 components to which a console processor can be connected and each component is assigned a specific port number. The eight components are the two diskette drives and one modem (ports 0 and 1), power (port 2), the IPPF/execution function and PSCF (port 3), the trace unit and counter (port 4), the first channel group (port 5), the second channel group (port 6), and the third channel group (port 7).

Port address 1 is used by one console processor to address the diskette drive of the other console processor. This function can be used, for example, when a customer engineer using the service support station wants to access the power logout records that are contained on the diskette drive of the operator station.

The configuration frame is used to select the ports that can be accessed by each console processor if the default assignment established during IMPL is not to be used. An actual connection between a component and a console processor is made when the console processor specifically addresses an assigned port via a microprogram instruction. Only one port can be addressed at a time by a console processor. A component must be powered up in order to be configured to a console processor. However, the diskette drives and power components are always configurable. For proper operation of the processor, each port should be assigned to one console processor or the other.

The state of an operating station determines the ports it can be assigned and the types of diskettes that can be mounted on its diskette drive as follows:

- <NO> - ports 0 and 2 and the system diskette.
- <NS> - ports 0 and 3 through 7 (as present and powered up) and the operational and processor diagnostic diskettes.
- <MO> - all ports except the one assigned to the <MS> station. Any port 4 through 7 can be assigned to the <MS> station. Any diskette except the director diagnostic diskette can be used.
- <MS> - Port 0 and the port (4 through 7) for the component to be serviced. Any diskettes except the processor diagnostic diskettes can be used.

NORMAL CONSOLE CONFIGURATION

The normal console configuration is shown in Figure 20.30.1. The normal configuration is established by default when a power-on IMPL is performed. The 3036 Console is in normal mode. One console station is assigned as the operator station and is in the <NO> state. The other is assigned as the service support station and is in the <NS> state.

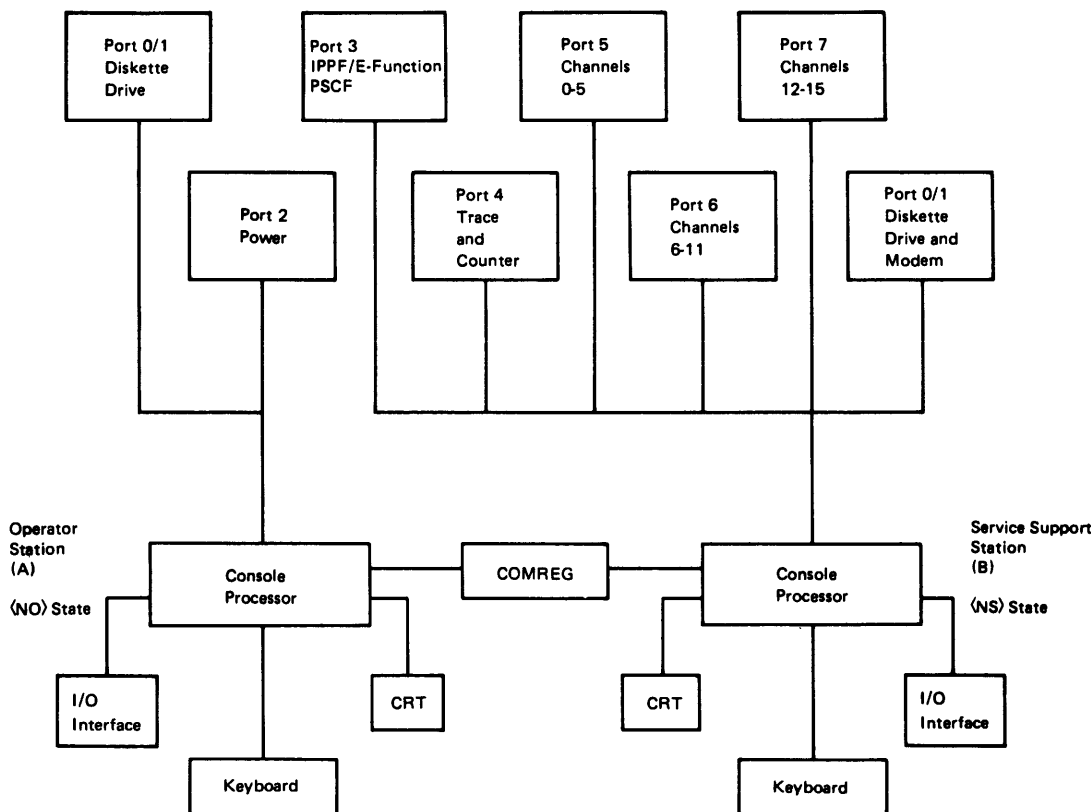


Figure 20.30.1. The normal 3036 Console configuration

The operator station (usually station A) can select its own diskette drive, which is addressed as port 0, and the power port. The

service support station can select its own diskette drive and (if it is station B) the modem as port 0 as well as ports 3 through 7.

When the normal configuration is in effect, the operator station is used for communication with the operating system, via the program frame, and can be used for power control if necessary (using the power control frame). All trace unit recordings, SSR records, and logout data from the 3033 Processor except power logouts are written to the diskette drive of the service support station. Status information for updating line 25 of certain display frames is also written to the service support station, which passes this data to the operator station for display on the program frame.

In the normal configuration, the service support station can be inactive (with any frame desired displayed, except power control), it can be used to perform monitoring or certain concurrent maintenance operations, or it can be used as an additional operator console using Multiple Console Support (MCS) in OS/VS.

The service support station can be used to perform the following concurrently with normal processing with the normal console configuration in effect:

- Set up modes of operation for the trace unit and enable/disable SSR using the mode control frame
- Display trace event records and associated logouts contained on its system diskette.
- Display processor logouts, director logouts, or system status recording records contained on its system diskette. Any power logouts contained on the mounted diskette can be displayed also.
- Monitor system operation using the operator frame or indicator frames associated with the instruction processor function or the directors
- Utilize the TP link frame to perform any one of the three possible remote maintenance functions: remote program, transmit logs/trace, or remote console. The remote program function can be used only to execute OLTs under OLTEP or to execute the data link software (DLS) program. The remote console function can be used only for remote monitoring of processor operations. In order to perform one of these remote maintenance functions, the service support station must be station B, since the modem is required.
- Access the diskette drive as a service record file using (1) an OLT to format and print trace data and logouts or (2) EREP to obtain the formatting information for MCH and CCH records

If the service support station is to be used as an operator console in the normal console configuration, program frame microcode must be loaded using the configuration frame, since service record file microcode is loaded as the default for the service support station in normal mode.

Remote Maintenance Operations

In order to utilize a remote maintenance function supported via the TP link frame, a communications line must be installed to connect the modem in the 3036 Console to the RETAIN/370 system. The RETAIN/370 system itself is connected to remote 3270 display devices. The remote 3270 displays are operated by customer engineer specialists.

Such a connection gives the remote customer engineer access to the 3033 Processor Complex for maintenance functions subject to the installation-established security. A local customer engineer and a remote customer engineer specialist can communicate via a message area in the TP link frame. Voice communication between the two can be accomplished using an additional telephone line between the installation and the remote specialist.

Only one of the three remote maintenance functions can be utilized at a time and they can be invoked using station B only. When one of these functions is active, the diskette drive of station B cannot be accessed as a service record file. In addition, while the remote program function is in execution, the program frame cannot be utilized on station B.

Before any one of the remote maintenance functions can be initiated, telecommunications must be activated using the control panel on the 3036 Console, as previously discussed. This activation enables two lines of customer information from the operational diskette to be transmitted to the remote specialist for security verification. If the information is valid, the remote specialist returns a ready message and telecommunications are established.

Remote Program Function. When the remote program function is selected, 2955 emulation is invoked. The remote customer engineer specialist can execute OLTs under OLTEP or OLTSEP, run the System Test/370 program, execute the diagnostic monitor program, or utilize the data link software program for software debugging. The remote specialist communicates only with the executing program (not with station B). Only the OLTEP and DLS programs can be executed concurrently with normal processing.

Transmit Logs/Trace Function. This function enables a remote specialist to request transmittal to RETAIN/370 of all or selected logout and trace records contained on the operational diskette. The remote specialist can obtain this data from RETAIN/370 for analysis using his display console.

Remote Console Function. This function can be invoked when the local customer engineer determines that the help of a remote specialist is required to locate a malfunction. This function enables the remote specialist to use a 3270 display to perform most of the same functions on the 3033 Processor as can be performed by the local customer engineer using the service support station. That is, the remote specialist can select a frame, request a function, and receive the results of the execution of the function at his 3270 display. The local and remote customer engineers can communicate using the TP link frame. The transmit logs/trace function can also be utilized when the remote console function is invoked (if the transmit logs/trace function is allowed).

MAINTENANCE MODE CONSOLE CONFIGURATIONS

A maintenance mode console configuration is established when the operator selects the "enter maintenance mode" command using the configuration frame. One console station is assigned as both the operator and service support station (in <MO> state) while the other console station is in the <MS> state and can be used by the customer engineer to perform certain maintenance operations concurrently with normal processing. Maintenance mode is also used when one of the operating stations has an inoperable component and when the processor is to be dedicated to maintenance operations.

A configuration in which both the operator and service support function are assigned to one console station is sometimes referred to as a degraded mode configuration because it can result in a slight degradation in response to the operator using the combined console station if an exceptional condition, such as a channel or processor error, occurs simultaneously with operator responses or requests. The performance of the processor itself is not affected by a maintenance mode console configuration.

Figure 20.30.2 shows a sample maintenance mode console configuration in which one console station is to be used for certain maintenance functions. The combined operator/service support station is in the <MO> state while the console to be used for maintenance is in the <MS> state.

All ports are assigned to the combined operator/service support station except the diskette drive of the other console station. All logouts, trace data, and system status recording records are written to the system diskette of the combined operator/service support station. All the functions that can be performed by the operator station and the service support station in the normal configuration, except for access to the diskette drive as a service record file and all remote maintenance functions, can be performed using the combined operator/service support station.

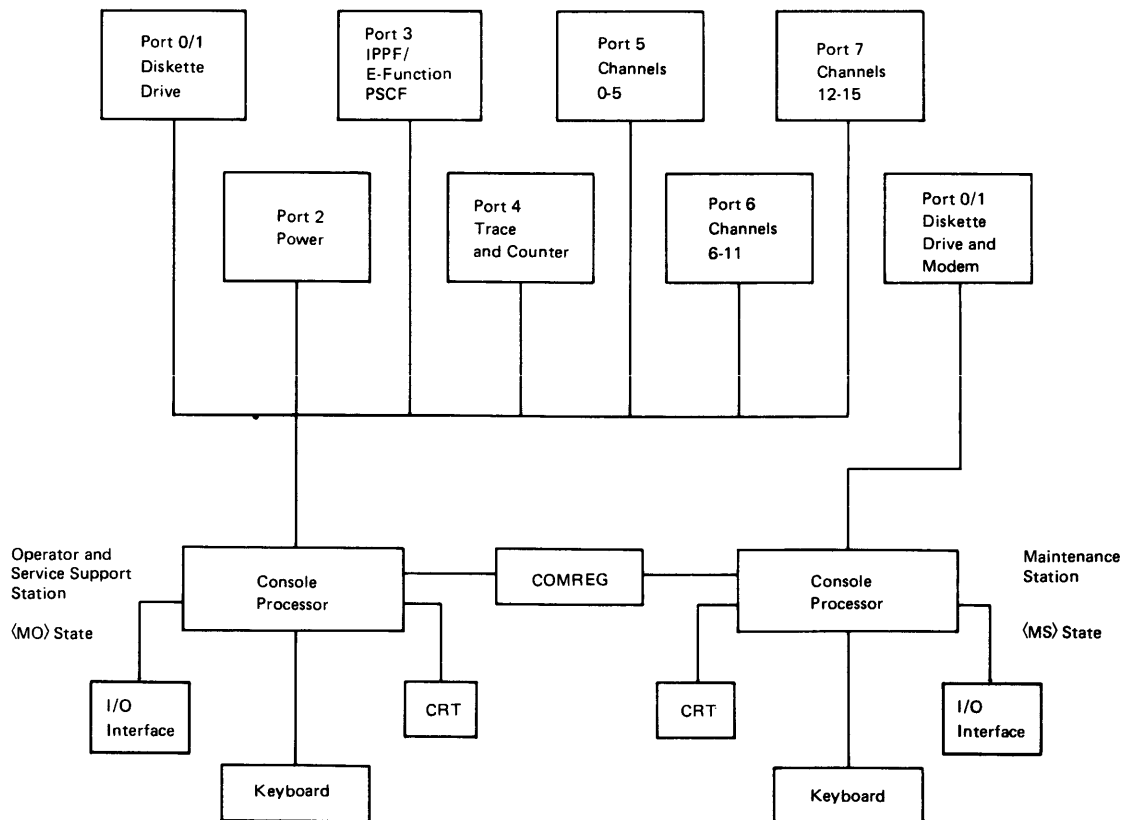


Figure 20.30.2. A maintenance mode 3036 Console configuration

When the 3036 Console is in the maintenance mode configuration shown in Figure 20.30.2, the station in the <MS> state can be used to perform the following concurrently with normal processing:

- Display processor, director, or any power logouts contained on the system diskette of this console station
- Display trace events and associated logouts contained on the system diskette of this console station
- Display SSR records on the system diskette of this console station
- Use the TP link frame to transmit logs and trace data from the system diskette of this console station to RETAIN/370, only if this station is station B
- Run console diagnostics on this console station, if it is malfunctioning

Maintenance operations can be performed on a malfunctioning director concurrently with normal operations using a maintenance mode configuration. To do this, the malfunctioning director port must be assigned to the console station that is in the <MS> state, as shown in Figure 20.30.3.

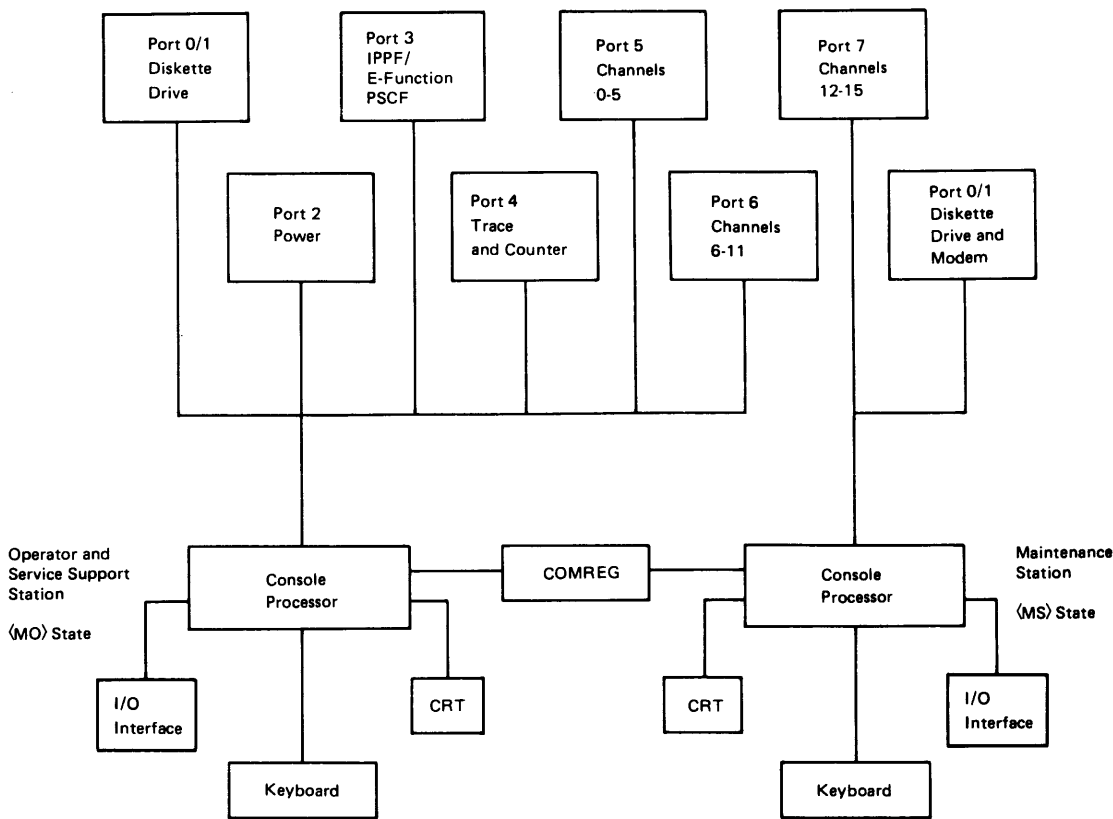


Figure 20.30.3. Director servicing using a maintenance mode console configuration

When the maintenance mode configuration shown in Figure 20.30.3 is in effect, the station in the <MS> state can be used to (1) execute director microdiagnostics contained on a diskette on the malfunctioning director, (2) exercise the malfunctioning director using manual control and indicator frames (or manual controls on the 3033 Processor), or (3) use the remote console function to enable a remote specialist to

exercise the malfunctioning director (only if station A is the combined operator/service support station).

If a director malfunction occurs during normal processing operations and the director is to be serviced, the configuration shown in Figure 20.30.3 can be established from the normal console configuration without a re-IMPL or re-IPL. Once the malfunctioning channel group is varied offline and processor operations are stopped, station B and the configuration frame can be used to set up the maintenance mode configuration in which station A is configured as the combined operator/service support station (in state <MO>) and console B is placed in <MS> mode. Console A is then used to assign the malfunctioning channel group port to station B.

Normal processor operations are then resumed using station A. The director diagnostics diskette must then be mounted for station B and IMPLed. When diagnostic execution is completed, the system diskette should be remounted for station B and IMPLed. Station A should be used to reconfigure the 3036 Console to normal mode.

When the 3033 Processor Complex is dedicated to servicing operations, the two console stations can be used to service two directors, a processor component and a director, a processor component and console station, or a director and a console station concurrently. Remote maintenance functions can be utilized for the component being serviced by station B. When the processor is powered off, diagnostics can be executed on both console stations at the same time. When the power diagnostics are executed using one station, the other station is inoperable.

CONSOLE RECOVERY FACILITIES

Certain recovery procedures occur automatically when a failure occurs on a console station. First, when normal mode is in effect for the 3036 Console, the two console stations monitor each other via their communication register. If one station determines that the other station has not responded to communication for 6 seconds, the functional station automatically enters the <MO> state and directs the failing station to IMPL. If the failing station can IMPL, it enters the <MS> state and displays a message to indicate the other operating station has assumed all console functions. If the station cannot IMPL (security switch is in the wrong position, for example) the other station assumes all functions but the message does not appear on the failing station.

Second, when an IMPL of a console station is not successful (error IMPL condition), and the diskette mounted for the station is not the console and power diagnostics diskette, an automatic re-IMPL is initiated. If the re-IMPL is successful, operations continue normally. If the re-IMPL is not successful, another re-IMPL is initiated. If 32 unsuccessful re-IMPLs are performed consecutively, a message is issued to indicate the fact and the operator should try to IMPL the station via the IMPL pushbutton. This error IMPL procedure is also performed automatically whenever a control storage error occurs for a console station during processing.

Third, when an IMPL of a console station is performed during processor operation with the 3036 Console in normal mode and the system diskette is mounted (allowing the basic console diagnostics to be executed) a message is issued if an error is found. The other operating station then automatically enters the <MO> state and assumes all console functions.

20:35 SERVICEABILITY FEATURES

The design and functions of the 3036 Console incorporate certain serviceability features of other large scale System/370 processors, such as the 3158 and 3168, and provide additional serviceability functions not available for these other processors. For example, while certain concurrent maintenance operations can be performed on 3158 and 3168 Processing Units, neither offers the operational advantages or range of concurrent procedures provided by the 3036 Console.

Specifically, a display console is not available for use with the service processor of a 3168-3 Processing Unit (and there is no service processor in a 3168-1), and in a Model 158 configuration there is only one display, which must be shared for both normal and service procedures during concurrent maintenance operations. In addition, more types of maintenance operations can be performed concurrently with normal processing in a 3033 Processor (logout data contained on the diskette drive can be displayed without interfering with other processing and channels can be serviced concurrently with normal processing). See Section 60 for detailed differences between maintenance operations for the 3033 Processor and 3168 Processing Unit.

The serviceability of the 3033 Processor is also aided by the voltage monitoring capability implemented in the 3036 Console and the fault location techniques that can be implemented as a result of 3036 Console design.

VOLTAGE MONITORING

Power sequencing in a 3033 Processor Complex is microcode controlled by the 3036 Console, as described previously, rather than by mechanical means (relays and switches). This design enables microprogram-controlled voltage monitoring to be implemented.

When an overvoltage or overcurrent condition other than that which causes a circuit breaker to trip occurs, a power fault interruption condition is presented to the console processor to which the power port is assigned (normally the operator station). The console processor then tests for a continuation of the overvoltage or overcurrent condition before it determines a power-off sequence should be initiated. Undervoltage and undercurrent conditions are also monitored.

When a power fault interruption is processed by the console processor, a power logout record is written to the diskette drive of the operator station. This logout indicates the date and time of day, identifies the power fault location, and describes the power fault condition. These logout records can be displayed.

This monitoring scheme enables transient power faults to be detected and located (which can be difficult when a circuit breaker is not tripped) and permits a faster reaction to a power fault.

FAULT LOCATING TESTS

For the 3168 Processing Unit, writable control storage (WCS) has to be loaded before a building block approach of microcoded fault locating testing can be initiated. However, a sizable portion of the processor must be operating correctly before WCS can be loaded. This results in a larger dependence on programmed diagnostics (which take longer to execute than microdiagnostics because of coding to clean up error conditions created for testing purposes) than is necessary for the 3033 Processor.

In a 3033 Processor Complex, the 3036 Console controls the loading and execution of microdiagnostics for the 3033 Processor, which are contained on the processor microdiagnostics diskettes. A building block approach of fault location testing is utilized from the start of testing. First, console microdiagnostics contained on the processor microdiagnostics diskette are executed on the console station being used, to ensure the correct operation of the console processor. If a failure occurs that prevents further testing using one console station, a switch to the other console station can be made.

Once console microdiagnostics have been successfully executed, processor beginning tests are loaded into control storage of the console processor. These microdiagnostics test the communication paths between the 3033 Processor and 3036 Console and test the 3033 Processor for its ability to execute additional microdiagnostics.

When the processor beginning tests execute successfully, console-controlled processor microdiagnostics can be loaded into control storage of the console processor and control storage of the execution function in the 3033 Processor from the processor microdiagnostics diskette. The microdiagnostics in the console processor initiate, monitor, analyze results of, and control execution of the microdiagnostics contained in the execution function, which test the processor and include the testing of communication lines between the processor and channels.

The console-controlled channel microdiagnostics contained on the director microdiagnostics diskette can then be loaded into director control storage to test director and channel controls under control of the console station. Once all of these microdiagnostics have executed successfully, programmed processor functional diagnostics can be executed under control of the diagnostic monitor program, if necessary.

SECTION 30: THE 3033 MULTIPROCESSOR COMPLEX

30:05 GENERAL DESCRIPTION

INTRODUCTION

The 3033 Multiprocessor Complex offers large-system users the advantages of shared storage multiprocessing. A 3033 Multiprocessor Complex configuration, which is similar in architectural design to a Model 168 Multiprocessing System, contains two interconnected 3033 Processors that contain multiprocessing hardware.

The two processors in a 3033 Multiprocessor Complex can share their processor storage and operate under the control of a single operating system that is resident in the shared processor storage. One input queue and one task queue can be maintained for the configuration and both processors can be used to process each task (but not simultaneously).

IBM-supplied programming systems support of shared storage multiprocessing for the 3033 will be provided in OS/VS2 MVS. The 3033 Multiprocessor Complex is upward compatible with Model 168 and Model 158 Multiprocessing Systems. Thus, user-written processing programs that operate under MVS in a Model 168 or 158 Multiprocessing System (or uniprocessor configuration) can operate in a 3033 Multiprocessor Complex without modification, subject to the constraints listed in Section 10:05.

The 3033 Multiprocessor Complex offers price performance advantages over Model 168 Multiprocessing Systems in addition to the availability and serviceability advantages of a uniprocessor 3033 configuration over a uniprocessor Model 168 configuration. It also provides a growth path for 3033 Attached Processor Complex installations, since it offers additional channels and availability features.

The price performance of the 3033 Multiprocessor Complex makes shared storage multiprocessing available to large-system users who desire the advantages provided by this type of configuration but who previously could not justify the cost. Because 3033 Multiprocessor Complex configurations are upward compatible with uniprocessor System/360 and System/370 processors, nondisruptive growth is provided for installations with multiple systems (Models 65, 155, 158, 165, and 168, for example) that require a shared storage multiprocessing environment.

The 3033 Multiprocessor Complex can be part of a wide variety of multiprocessing configurations. The breadth and flexibility of the multiprocessing configurations supported for System/370 enable an installation to combine multiple processors such that the particular advantages offered by different types of multiprocessing configurations can be obtained as desired and ease of growth is assured.

GENERAL DEFINITION OF MULTIPROCESSING

A multiprocessing configuration is one in which two or more processors are interconnected and execute two or more tasks simultaneously, one in each processor. Multiprocessing is a logical extension of multiprocessing, in which two or more tasks operate concurrently in a single processor. In a multiprocessing environment, one task executes at a time and only I/O operations for two or more tasks can operate simultaneously. In a multiprocessing environment, both I/O operations and instruction execution for two or more tasks in the same or different programs can occur simultaneously, with each task executing in a different processor.

The hardware connection of the processors in a multiprocessing configuration is the means by which the processors communicate with each other in order to coordinate the activity of the multiprocessing configuration. A multiprocessing configuration can be tightly or loosely coupled or can include a combination of both loosely and tightly coupled processors.

A tightly coupled multiprocessing configuration is one in which (1) the processors share access to all the processor storage available in each system, (2) processor-to-processor communication is accomplished via the storing of data in shared storage and via direct processor-to-processor signals (both program- and hardware-initiated), and (3) a single control program is used. The 3033 Multiprocessor Complex is, therefore, a tightly coupled multiprocessing configuration, as is the 3033 Attached Processor Complex.

A loosely coupled multiprocessing configuration is one in which (1) processors are coupled via shared access to direct access storage or via channel-to-channel connections, (2) each processor has its own control program, and (3) a single system scheduling and operational interface is optional.

Loosely coupled multiprocessing configurations for System/370 processors are supported by ASP (Asymmetric Multiprocessing System) Version 3, JES2 Multi-Access Spool support in OS/VS2 MVS, and JES3 support in OS/VS2 MVS. ASP Version 3 supports from 2 to 33 systems connected via Channel-to-Channel Adapters. One support system schedules the operation of up to 32 main systems.

The JES2 Multi-Access Spool facility supports from two to seven systems, each with its own control program resident, that share input and output work queues on direct access storage shared by the systems. JES3, a generally compatible extension of ASP, supports from two to eight systems connected via Channel-to-Channel Adapters. Each of the systems in a JES3 configuration can be a uniprocessor system or tightly coupled multiprocessing system (that is, 3033 Multiprocessor Complex, 3033 or 3031 Attached Processor Complex, Model 158 or 168 Multiprocessing System, or Model 158 or 168 Attached Processor System).

The objective of coupling multiple systems to form a multiprocessing configuration is to obtain a configuration that combines advantages of a single processor environment with those of an uncoupled multiple processor environment.

A single processor environment offers the following advantages:

- A single interface to the computing system for workload scheduling and operation of the system
- The ability to apply all the resources of the system to a given job step when necessary

The advantages provided by an uncoupled multiple processor configuration are:

- The capability of adding to the configuration in smaller increments, that is, the addition of a smaller processor rather than replacement of the existing processor with the next larger processor when additional computing power is required. The next larger processor may provide additional computing power far in excess of that required.
- More economical growth possibilities for installations with purchased systems

- Growth possibilities for large-scale installations that have the largest processor of the system already installed
- Enhancements to configuration availability (better probability that a system will be available for critical application processing), concurrent maintenance, and improved reliability (protection of critical jobs from failures in noncritical jobs by processing them in separate systems)

THE 3033 MULTIPROCESSOR COMPLEX CONFIGURATION

The 3033 Multiprocessor Complex, as shown in Figure 30.05.1, consists of the following: (1) two multiprocessor models of the 3033 Processor (any processor storage model) interconnected via the 3038 Multiprocessor Communication Unit, (2) two 3037 Power and Coolant Distribution Units with a multiprocessing feature installed, and (3) two 3036 Consoles. Two motor generator sets, one for each processor, are also required. The multiprocessor model of the 3033 Processor contains a multiprocessing function that is required for a tightly coupled configuration.

The physical dimensions of the multiprocessor models of the 3033 Processor, 3037 Power and Coolant Distribution Units, and 3036 Consoles used in a 3033 Multiprocessor Complex are the same as those in a uniprocessor configuration.

When multiprocessor mode is in effect and OS/VS2 MVS multiprocessing support is used, the two processors in a 3033 Multiprocessor Complex configuration share from 8M- to 16M-bytes of processor storage, I/O devices that have channel- or string-switching features installed on their control units, a single control program, and a single set of work (input and output) queues. The channels in each processor are normally dedicated to that processor.

The two processors in a 3033 multiprocessing configuration can be reconfigured using the 3036 Console to permit each processor complex to operate independently of the other, that is, in uniprocessor mode.

Processor Features

Uniprocessor models of the 3033 (including the no longer available Model U6) can be field-converted to multiprocessor models and connected to the 3038 unit. The multiprocessing feature for the 3037 is also field-installable. Thus, two 3033 Processor Complex configurations can be field-converted to a 3033 Multiprocessor Complex configuration.

The two 3033 Processors in a 3033 Multiprocessor Complex have the same standard and optional features as the processor in a 3033 Processor Complex with the same prerequisites and limitations. All standard and optional features can operate when multiprocessor mode is in effect. However, IBM-supplied programming systems support for the 3033 Multiprocessor Complex supports tightly coupled multiprocessing operations only for EC and DAT modes. Multiprocessing operations are not supported for BC mode.

OS/VS2 tightly coupled multiprocessing support does not require symmetric processor features and provides a processor affinity facility that can be used to ensure that a program requiring a feature present in only one processor is processed only by that processor. However, this affinity support is not required for a 3033 Multiprocessor Complex configuration because of the nature of the optional features available (Channel-to-Channel Adapter, Two-Byte Interface, and Extended Channels). The affinity facility can be used, if desired for performance reasons, to ensure that a program that accesses an I/O device attached to only one processor executes only in that processor.

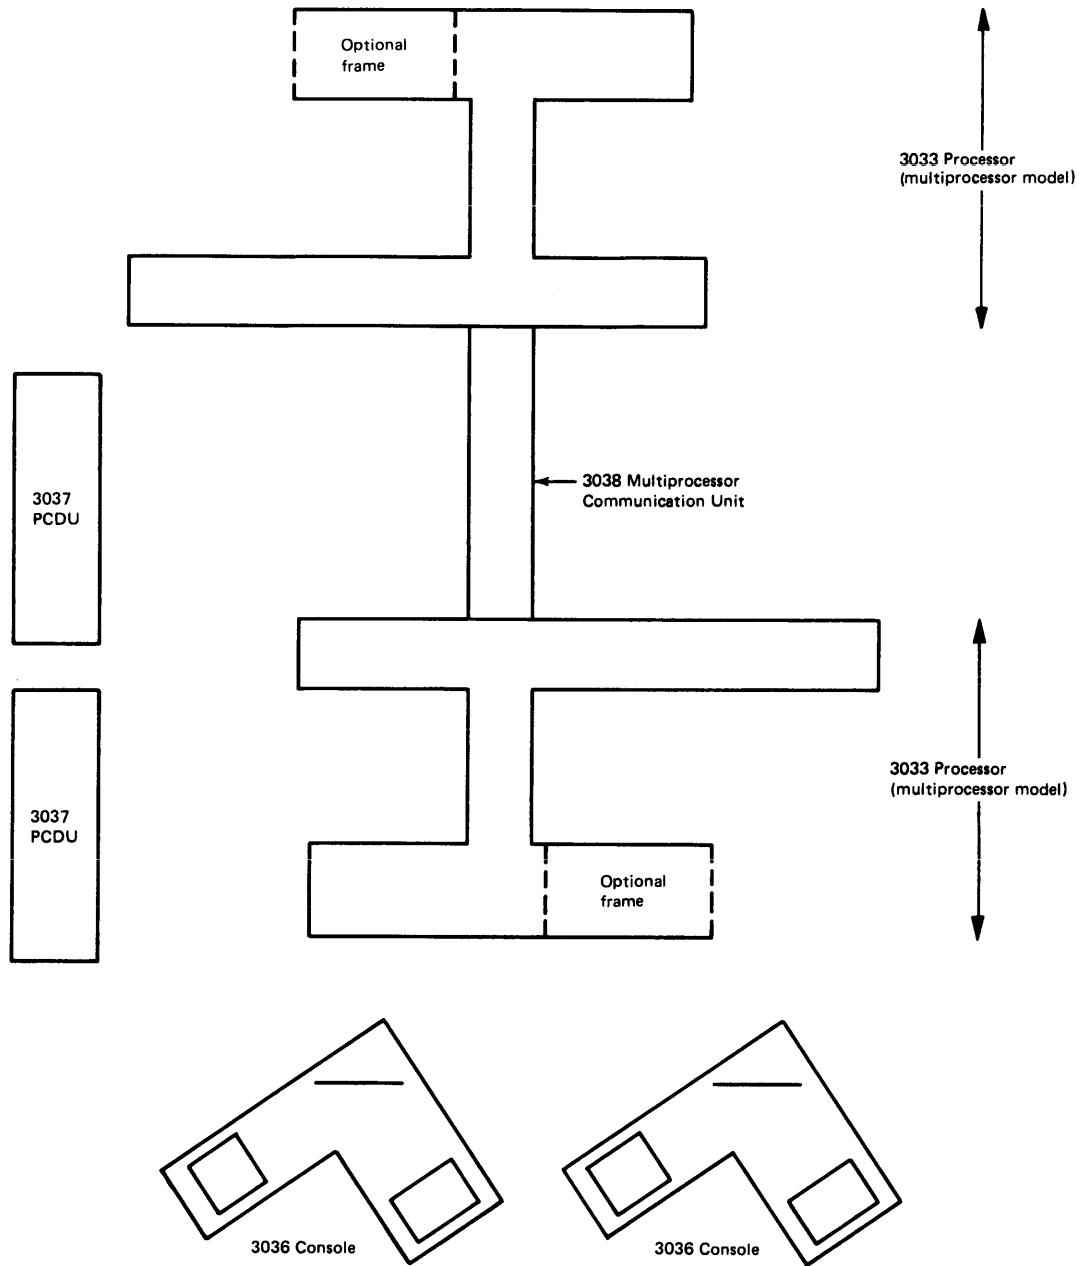


Figure 30.05.1. The 3033 Multiprocessor Complex

Processor Storage

The processor storage used in the two processors in a 3033 Multiprocessor Complex is the same eight-way, doubleword interleaved monolithic storage that is used in uniprocessor 3033 configurations. The storage sizes for multiprocessor models of the 3033 Processor are the same as for uniprocessor models and are designated as Models M4, M8, M12, and M16.

Asymmetric processor storage multiprocessing configurations (the connection of two 3033 Processors with different amounts of processor storage) are permitted for the 3033 Multiprocessor Complex. The total amount of processor storage in a 3033 multiprocessing configuration can vary from a minimum of 8M-bytes (two M4 models) to a maximum of 32M-bytes (two M16 models) in 4M-byte increments.

However, the maximum amount of processor storage that can be addressed during multiprocessor mode operations is 16M-bytes. That is, when more than 16M-bytes is installed in the configuration, only 16M-bytes can be enabled when multiprocessor mode is in effect. During uniprocessor mode operations, each 3033 Processor can address a maximum of 16M-bytes.

The total processor storage contained in each processor in a 3033 Multiprocessor Complex is logically divided into 2M-byte units, called storage bands, in order to implement floating storage addressing. Floating storage addressing is implemented so that the physical address range assigned to each band of processor storage can be varied as needed. The address range of each 2M-byte unit of processor storage is specified using a configuration (C2) frame of the 3036 Console.

Since processor storage is eight-way interleaved in a 3033 Multiprocessor Complex, the total processor storage available in the multiprocessing configuration is divided into eight logical storage elements addressed 0 through 7, each of which can be accessed concurrently. Each logical storage element consists of a portion of the processor storage physically contained in each processor. For an 8M-byte configuration, for example, the first logical storage element consists of the first 512K in each processor. The second consists of the second 512K in each processor, etc.

Nonsimultaneous requests from the two 3033 Processors to a nonbusy logical storage element are handled on a first-come, first-served basis. When both processors simultaneously request access to processor storage, storage priority is granted on a priority-weighted demand basis. That is, a given processor retains access priority to storage and can access storage on successive cycles (rather than every alternate cycle) as long as the processor can sustain storage requests of a higher priority classification than the other processor. Thus, the processor that has storage priority retains this priority until the processor without storage priority has a higher priority request for storage than the processor currently holding storage priority. In this situation, storage priority is switched from one processor to the other.

When the multiprocessing feature is present in a 3033 Processor, requests from a processor to its own physical processor storage (local requests) and to processor storage contained in the other processor (remote requests) are channeled to storage via the 3038 Multiprocessor Communication Unit. Therefore, the same amount of time is required for a remote request as for a local request when multiprocessor mode is in effect and more time is required to access processor storage in a 3033 Multiprocessor Complex than in a 3033 Processor Complex.

Channels

The same number of channels permitted in a uniprocessor model of the 3033 (12 or 16) can be installed in each processor in a 3033 Multiprocessor Complex. Thus, the 3033 Multiprocessor Complex can have a total of 24, 28, or 32 channels. Ideally, each processor should have the same number and types of channels so that at least two channel paths to I/O devices, one from each processor, can be made available via installation of programmable channel- or string-switching features. Asymmetric channel configurations are supported so that a malfunctioning channel can be logically removed from the operational system configuration during processing.

Normally the channels in a 3033 Processor are dedicated to that processor and cannot be accessed by the other processor in the multiprocessing configuration. However, channel set switching hardware is included in the two processors in a 3033 Multiprocessor Complex as part of the multiprocessing design. Channel set switching, which is not provided in Model 158 and 168 Multiprocessing Systems, is designed to aid the recovery provided by the OS/VS2 multiprocessing control program after one processor fails and must be logically removed from the operational multiprocessing configuration. Channel set switching can be activated by programming or the 3036 Console (C2 configuration frame).

The Channel Reconfiguration Hardware (CRH) function in OS/VS2 MVS for the 3033 Multiprocessor Complex will use channel set switching to permit channel controls in the failing processor to be shared alternately with channel controls in the functional processor under program control without a re-IPL after an uncorrectable processor failure occurs. The data paths of the failing processor must be functional and power must be on in the units through which the data lines pass, that is, certain 3033 Processor frames, the 3038 unit, and the associated 3036 Console and 3037 Power and Coolant Distribution Unit. Activation of channel set-switching permits the functional processor to receive and process I/O interruptions from, and initiate I/O operations to, I/O devices attached only to the failing processor.

I/O Devices

Any I/O device that can be attached to the channels in a 3033 Processor in a 3033 Processor Complex can be attached to the channels in the processors in a 3033 Multiprocessor Complex. While OS/VS2 multiprocessing support does not require symmetry for all I/O devices, for maximum availability the I/O device configuration should be as symmetric as possible. Ideally, the same I/O device configuration should be attached to each processor and, where it is available, a programmable channel-switching or direct access device string-switching feature should be installed on each control unit to provide each processor with access to the device.

A maximum of four channel paths per processor to a given control unit is supported by OS/VS2 multiprocessing support. (Note that only two channel paths to a device are supported for uniprocessor configurations in OS/VS2 MVS.) A control unit can be attached to: only one channel in only one processor, two channels in only one processor, one channel in each processor, or two channels in each processor. The latter two configurations are preferable for tightly coupled multiprocessing configurations. I/O devices connected to control units with a channel- or string-switching feature must have the same I/O address (channel, control unit, and device) in each processor.

Note that the total number of I/O devices present in a 3033 Multiprocessor Complex configuration cannot be greater than the maximum number of I/O devices present in a uniprocessor configuration because OS/VS2 MVS supports the same maximum number of UCBs (unit control

blocks) for I/O devices in its multiprocessor support as in its uniprocessor support.

String-switching features are available for 3330-series, 3340/3344, and 3350 disk storage and for the 3330 strings in a 3850 Mass Storage System. The following are some of the more frequently used I/O devices for the 3033 Processor that have a programmable two-channel switch available for their control unit:

- 1403 Printer
- 2540 Card Read Punch
- 2400-series Magnetic Tape Units
- 3400-series Magnetic Tape Units
- 2314/2319 Disk Storage
- 2305 Model 1 and 2 Disk Storage
- 3330-series Disk Storage
- 3340 and 3344 Disk Storage
- 3350 Disk Storage
- 3800 Printing Subsystem
- 3850 Mass Storage System

The Channel Adapter Type 3 feature can be installed on a 3705 Communications Controller to permit it to be switched between the two processors in a 3033 tightly coupled multiprocessing configuration under program control. This channel adapter is functionally equivalent to the programmable two-channel switch available for certain direct access devices (except for RESERVE/RELEASE functions). The 3704 Communications Controller has a nonprogrammable two-channel switch that permits it to be manually switched between two processors.

For other I/O devices, the 2911 Manual Switch or 2914 Switching Unit can be installed to provide a manual switching capability for their control units when the device is not present in the I/O configuration of both processors. Manual switching using the 2911 is normally done by the operator when the processors are not operating, such as before an IPL. This is also true for a 2914 unless additional RPOs are installed on the 2914 that enable switching to be done dynamically during processor operation (switching becomes effective the next time the channel interface becomes inactive).

The channel configuration rules suggested for 3033 uniprocessor configurations (see Section 10:20) should also be observed for each processor in a 3033 multiprocessing configuration. In addition, channel and I/O device symmetry should be planned.

3038 Multiprocessor Communication Unit

The 3038 Multiprocessor Communication Unit physically connects the two processors in a 3033 Multiprocessor Complex and provides a communication path between the two processors and between the two 3036 Consoles. Functionally, the 3038 is divided in half. Each half is associated with the processor it is attached to and receives its power and water cooling from the 3037 of its associated processor. Each 3038 half can be powered up and down separately from the other half and contains an oscillator for timing its associated processor.

The processor storage lines from processor storage in each processor and the port 4 (trace function) lines for each 3036 Console are routed through the 3038 unit. Each half of the 3038 also has a set of registers that is used for communication between the two 3036 Consoles. Communication between the two 3036 Consoles is required, for example, to pass status information between the two processors and to execute certain orders of the SIGNAL PROCESSOR instruction.

3036 Consoles

The two 3036 Consoles in a 3033 Multiprocessor Complex are physically identical to the 3036 Console used in a 3033 Processor Complex. No changes are made to the control panel or the two operating stations. However, the operational diskettes for a 3033 multiprocessing configuration contain processor and console microcode that is designed to support multiprocessing and interprocessor communication capabilities.

For multiprocessor mode operations, at least one operating station in each 3036 Console must be functional. This is necessary for power control for each processor and 3036 Console-to-3036 Console communication, which is accomplished via the service support stations.

In the normal console configuration for multiprocessor mode operations, one station in each 3036 Console is designated as an operator station while the other is designated as a service support station. One operator station is used as the primary operating system console while the other can be defined as an alternate or additional console. The port assignments for the two stations in each 3036 Console are the same as the port assignments for the two stations in the normal console configuration for 3033 uniprocessor configurations, as shown in Figure 30.05.2.

For the normal configuration, the service support station of each 3036 Console can be utilized to perform the same service operations for its associated processor concurrently with normal processing as in a 3033 uniprocessor configuration, as described in Section 20:30 under "Normal Console Configuration". All logouts from a processor are written to the diskette drives of its associated 3036 Console.

If the operating station being used as the primary operating system console malfunctions, the operator station of the other 3036 Console can be utilized as the primary console. The service support station of the 3036 Console with the malfunctioning station can be assigned as a combined operator/service support station (in <MO> state) and console diagnostics can be run on the malfunctioning console station.

If both operating stations of a 3036 Console are malfunctioning, the associated processor cannot be utilized and multiprocessor mode of operation is not possible. However, uniprocessor mode operations using the other processor are possible.

The display frames provided for the 3036 Consoles in a 3033 Multiprocessor Complex provide the same functions as for a 3033 Processor Complex. Certain frames have been modified and additional frames are provided to support functions required in a tightly coupled multiprocessing configuration.

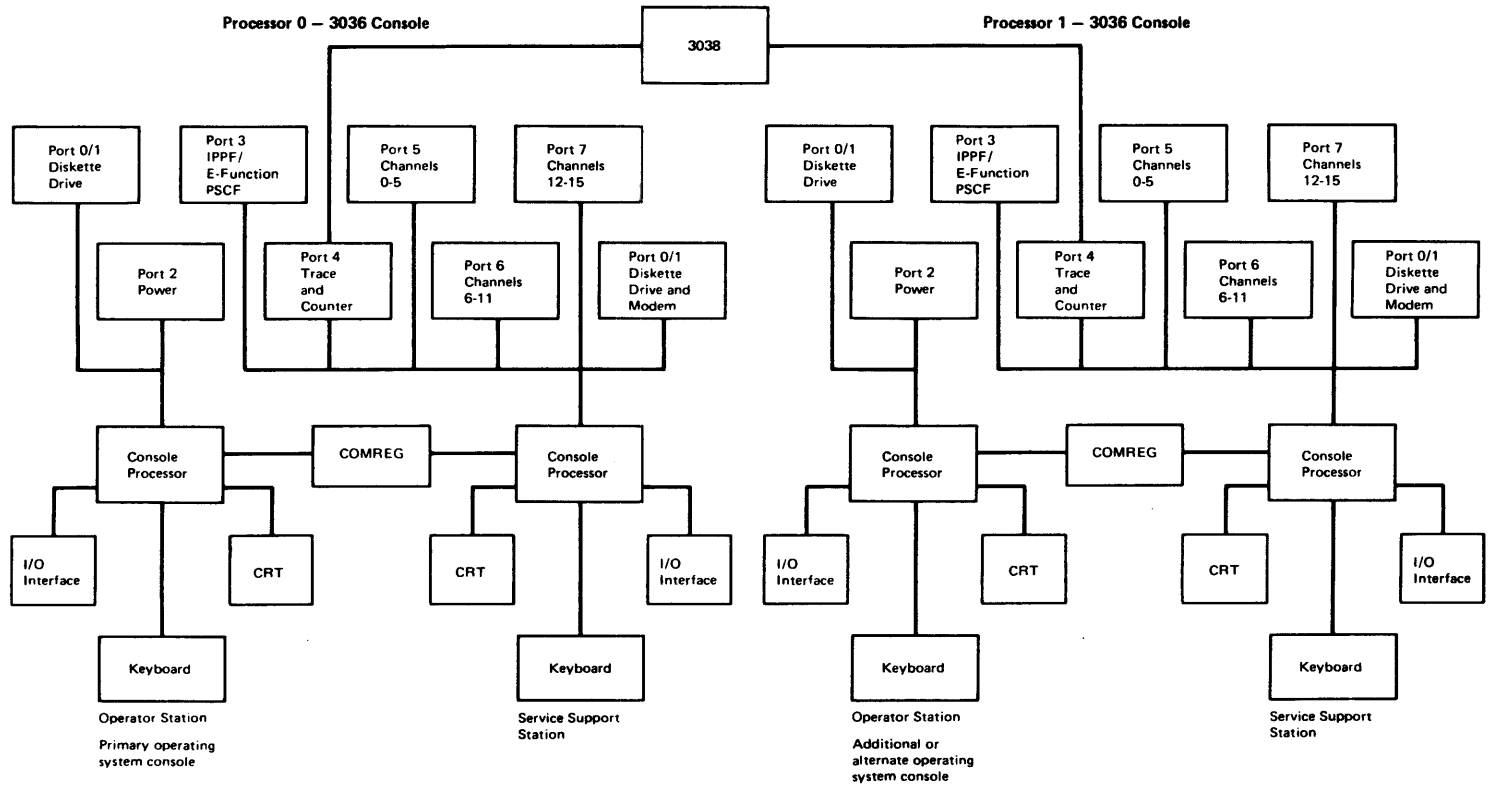


Figure 30.05.2. The normal console configuration for the 3033 Multiprocessor Complex

A second configuration (C2) frame is provided for a 3033 Multiprocessor Complex to perform configuration functions that in a Model 168 Multiprocessing System are accomplished using the configuration control panel on the 3068 Multisystem Unit. The C2 configuration frame is the only frame that permits an operator to establish configuration information for both processors in a 3033 Multiprocessor Complex using either 3036 Console (the one designated as the primary operator console). For other configuring for a processor (such as UCW assignments, console designations, powering, etc.), the 3036 Console associated with that processor must be used.

The C2 configuration frame is used to establish the following:

- Processor storage assignment and interleaving. Each 2M-byte processor storage band in the configuration has the capability of being enabled for access by both processors, disabled for access by both processors, or enabled for access by only one of the processors. An address range must be assigned to each enabled band. A maximum of 16M-bytes can be enabled during multiprocessor mode operations. Eight-way doubleword interleaving or serial operation can also be established, as for a uniprocessor configuration.

For multiprocessor mode operations, each processor storage band that is enabled to one processor must also be enabled to the other processor. This assignment permits both processors to access all of the enabled processor storage in the configuration. For uniprocessor mode operations, no processor storage band can be enabled for access by both processors; however, one processor can be allocated more processor storage than the other (that is, storage allocation need not be symmetrical). When uniprocessor mode is in effect and one or more bands of processor storage that are physically contained in one processor are enabled for access only by the other processor, processor storage is said to be cross-configured.

Each enabled processor storage band must have a 2M-byte address range assigned. Any of the eight possible 2M-byte ranges (0 to 2M through 14M to 16M) can be assigned to a storage band as long as the same address range is not assigned to more than one enabled band for multiprocessor mode operations or to two bands allocated to the same processor for uniprocessor mode operations. The address range 0 to 2M must be assigned to one enabled band for multiprocessor mode operations and to one enabled band for each processor for uniprocessor mode operations; otherwise, a successful IPL cannot be performed.

The address ranges assigned to the storage enabled to a given processor for uniprocessor mode operations or to shared storage for multiprocessor mode operations need not be contiguous when an OS/VS2 MVS multiprocessing operating system is used. In addition, address ranges higher than the total amount of processor storage enabled or present in the multiprocessing configuration can be assigned.

For example, address ranges 0 to 6M and 8M to 10M could be assigned to a 3033 multiprocessing configuration that contained 8M-bytes of processor storage. The addresses between 6M and 8M are marked unavailable by the OS/VS2 multiprocessing control program. The processor storage in each processor must be assigned contiguous address ranges for uniprocessor mode operations when an operating system other than OS/VS2 MVS with multiprocessing support is used.

- Uniprocessor or multiprocessor mode of operation. When multiprocessor mode is in effect, interprocessor communication hardware is enabled (SIGNAL PROCESSOR interface, high-speed buffer intercommunication, time-of-day clock security switch broadcasting, time-of-day clock synchronization check, malfunction alert signal,

and the broadcast of reset functions). These facilities are discussed in the remainder of this section.

When uniprocessor mode is in effect, the interprocessor signals are deactivated and each processor can operate independently of the other, with that portion of available processor storage and the I/O devices enabled to it by the operator.

- Oscillator assignments. The system oscillator for each processor in a 3033 Multiprocessor Complex is contained in its associated half of the 3038 Multiprocessor Communication Unit and the system oscillator in each processor is disabled. This is done so that a processor need not be operational in order for its processor storage to be accessed by the other processor. Each processor still contains its own time-of-day clock oscillator.

For multiprocessor mode operations, both processors use the same system oscillator and time-of-day clock oscillator so that operations are synchronized. The system oscillator not being used is phase locked to the controlling system oscillator so that control can be switched from the latter to the former without quiescing system operation. For uniprocessor mode, both may share one set of oscillators or each processor can use its own two oscillators. When processor storage is cross-configured, one set of oscillators must be used by both processors.

- Status of channel set switching. When the normal setting is selected, channel set switching is not active and each processor controls the operation of its own channels. When the reverse setting is selected, channel set switching is activated and control of the channels in each processor is switched to the other processor.
- Console(s) to be used for configuration control. The C2 configuration frame can be used to establish control for configuration changes. Changes can be made from both 3036 Consoles using the C2 configuration frame (each 3036 is a master console) or only one 3036 Console (the master). The 3036 Console that cannot make changes (designated the slave console) can only display the configuration in effect using the C2 configuration frame.

The ability to enable and disable the I/O interface to each processor for control units with the Remote Switch attachment installed (provided via the configuration control panel for a Model 158 or 168 Multiprocessing System) is not provided via the C2 configuration frame for a 3033 Multiprocessor Complex. This enabling and disabling must be done via the I/O interface switches on the control units.

Some possible valid configurations for the 3033 Multiprocessor Complex are the following:

- A full multiprocessing configuration, that is, multiprocessing mode in effect with all available processor storage enabled to both processors. This is the normal configuration.
- A dual uniprocessor configuration, that is, uniprocessor mode in effect with the processor storage of each processor enabled only to that processor (storage is nonshared and not cross-configured)
- Uniprocessor mode in effect with processor storage of only one of the 3033 Processors available. The available processor storage is divided between the two processors (each band enabled to only one processor). A minimum of 4M-bytes of processor storage must be available in one processor for this configuration.

- Uniprocessor mode in effect with 2M-bytes of storage allocated to one processor only (a maintenance or test system) and all the remaining available storage allocated only to the other processor.
- Multiprocessor mode in effect with the processor storage of only one processor enabled to both processors. A minimum of 2M-bytes of processor storage must be available in either processor for this configuration.
- Uniprocessor mode in effect with all available storage from both processors enabled to one processor. The other processor is not part of the production processing configuration, and diagnostics that do not require processor storage can be executed on it.

Note that when multiprocessor mode is in effect, a malfunction in a processor may cause the processor to be logically removed from the operational configuration by the OS/VS2 MVS multiprocessing control program. Multiprocessor mode is still in effect and can remain effective as long as nothing is to be executed in the malfunctioning processor. For diagnosis of the failure or repair, a maintenance subsystem must be configured. At this time, uniprocessor mode must be established.

When a valid configuration is entered via the C2 configuration frame, each processor is forced to a quiesced state at the end of the execution of its current instruction so that the specified configuration can be made effective. The configuration currently in effect is maintained in storage within the 3036 Console.

When multiprocessor mode is in effect, the specified configuration becomes effective in both processors simultaneously. When uniprocessor mode is in effect, a change becomes effective in one processor and does not affect the other processor.

As for a Model 168 Multiprocessing System, processor operations need not be quiesced prior to any alteration of the physical configuration using the C2 configuration frame in a 3033 Multiprocessor Complex. Whenever a configuration change is made, all the entries in one or both (as appropriate) high-speed buffers and translation lookaside buffers are invalidated.

The following communication facilities are provided for multiprocessor mode operations: broadcast of reset functions from one processor to the other, broadcast of time-of-day clock security switch enable set condition, and 3036 Console-to-3036 Console communication.

When a system reset (with or without the system clear) or load (IPL with or without a system clear) is performed for one system using the operator frame, a corresponding function is also automatically sent (broadcast) to the other system when multiprocessing mode is in effect, as shown below.

<u>Function Selected by Operator on Local Processor</u>	<u>Function Performed on Local Processor</u>	<u>Function Broadcast to Remote Processor</u>
System reset (normal)	Program reset	Program reset
System reset (clear)	System clear	Initial program reset
Load (normal)	Program reset	Program reset
Load (clear)	System clear	Initial program reset

The only other communication function provided in the 3036 Console for multiprocessor mode operations is for the time-of-day clock. When the time-of-day clock security switch for either system is held in the enable set position, this setting is propagated to the other processor so that the time-of-day clock in each system is enabled for setting.

Power Control

Power control for a 3033 Multiprocessor Complex is designed to allow certain components within one system to be powered on and off separately from other components of the given system without preventing operation of the other system.

Hardware and programming systems support for the 3033 Multiprocessor Complex are designed to enable a properly isolated system component to be powered down without the necessity of first quiescing system operations. This capability enables repair operations to be performed on a malfunctioning component while normal system operations continue, using the other system and functional components of the system with the malfunctioning component. When the malfunctioning component is repaired, it can be powered on and returned to the functional configuration, again without quiescing system operations.

Each of the system units that can be separately powered on and off has its own power control switch that can be set to the local or remote position. The switch must be set to local in order to perform separate power on and power off operations.

The following components within a given system in a 3033 multiprocessing configuration can be powered on and off separately from the other components in the same system:

- The processor storage in a 3033 Processor. Multiprocessor mode operations utilizing the IPPF, execution function, channels, and I/O devices of the 3033 Processor without storage powered up is possible.
- Each director/channel group in the processor. The balance of this processor, the switched I/O devices attached to the powered off channels, and I/O devices attached to other channels are available for multiprocessor mode operations.
- The IPPF, execution function, and PSCF of one processor. The processor storage of this processor is available for multiprocessor mode operations. The channels of this processor cannot also be utilized, since the PSCF must be powered on in order to utilize the channel set switching function.
- The half of the 3038 Multiprocessor Communication Unit associated with the system. The 3033 Processor and 3036 Console associated with the powered down half cannot be used.
- Individual control units and I/O devices attached to the channels of the system. The balance of the system is available for multiprocessor mode operations.

Note that when the 3037 unit for a system is powered down, the associated 3033, 3036, and associated 3038 half are also powered down. In addition, when the 3036 is powered down, it causes a powering down of its associated 3033 and 3038 half also.

For proper system operation, a component (processor, processor storage, channel, or I/O device) should be logically varied offline to the operating system before it is powered off. A component must be varied online after it is again powered on if it is to be part of the operational configuration.

ADVANTAGES OF TIGHTLY COUPLED MULTIPROCESSING CONFIGURATIONS

The major advantages of a tightly coupled multiprocessing configuration when compared with two uncoupled systems, each having half the resources of the total multiprocessing configuration, are:

- Improved availability
- Less complex operational requirements
- Improved resource utilization
- Operational flexibility
- Improved growth options
- Improved throughput possibilities

These advantages are made possible by hardware resource redundancy, extensive hardware reconfigurability (implemented both in hardware and the OS/VS2 MVS multiprocessing control program), tightly coupled hardware interconnection that permits the configuration to operate with one control program and one work queue, and the availability features that are basic to the design of the OS/VS2 MVS control program. In a tightly coupled multiprocessing environment, the most critical component is the control program, and processors are viewed as resources that are to be allocated to tasks just as are I/O devices, processor storage, and programs.

Availability

Availability as it relates to a data processing installation is usually described as the percentage of scheduled time the system or an application is capable of processing. A system is available when both its hardware and programming system are capable of processing jobs. An application is available when it is capable of performing processing for its end users.

A system is unavailable when, for example, a hardware or control program failure occurs and system recovery procedures are invoked, an operator error causes a failure and recovery is required, scheduled preventive maintenance is performed, engineering changes are applied to hardware, fixes are applied to programming systems, and diagnostics are performed to locate a hardware malfunction that prevents continued system operation.

The improved availability offered by a tightly coupled 3033 multiprocessing configuration is the result of hardware component redundancy, hardware component reconfigurability, and availability features implemented in OS/VS2 MVS uniprocessor and tightly coupled multiprocessor support. Since there are two processors and, in general, processor feature, channel, and I/O device symmetry as well as access to I/O devices by both processors, backup is usually available when a hardware component fails.

In addition, the reconfiguration capabilities supported by the OS/VS2 MVS multiprocessing control program enable any type of hardware resource (a processor, certain storage band, channel, or I/O device) to be logically removed from the operational multiprocessing configuration on an individual basis (within certain limitations) without the necessity of terminating system operations and performing a re-IPL. The rest of the multiprocessing system can then continue to function in degraded mode.

Availability is also enhanced by the capability of physically removing malfunctioning components from the configuration for deferred or concurrent maintenance without impacting the availability of the rest of the multiprocessing configuration. Once the operator issues a VARY OFFLINE command to indicate the component to be removed, the OS/VS2 multiprocessing control program issues a message as soon as it has finished using the component and has logically removed it from the operational system. Using the multiprocessing configuration frame of the 3036 Console, the operator can then physically remove the component from the configuration while normal processing continues, provided the proper operational procedure is followed.

In a configuration with two uncoupled systems, only unallocated I/O devices and redundant channel paths can be logically removed from the operational system (varied offline) without stopping the system. When a processor or a critical portion of processor storage is the malfunctioning component, all the hardware resources of the system are unavailable (except switched I/O devices) until repair has been completed.

While a configuration with two uncoupled systems does provide backup by offering component redundancy, a tightly coupled multiprocessing configuration provides other availability advantages. First, since there is a single work queue, instead of two independent job queues, in a tightly coupled multiprocessing configuration, recovery from a processor failure is faster because the switchover from one system to the other of jobs in progress and queued jobs is eliminated.

Second, when a failure occurs in a tightly coupled multiprocessing configuration, the hardware components of the failing system are available to assist in the recovery. In the case of a processor failure, the functional processor is available to perform the recovery. These capabilities are not provided in a configuration with two uncoupled systems.

A tightly coupled multiprocessing configuration also provides a better method of handling preplanned scheduled maintenance activities for installations that operate 24 hours a day, since a processor and its components can be physically removed from the multiprocessing configuration without the necessity of job cancellation and system restart.

Functions designed to increase system availability are basic to the structure of OS/VS2 MVS. For both uniprocessor and multiprocessor systems, OS/VS2 MVS is designed to attempt to reduce the frequency of errors that occur as the result of programming and to reduce the impact of both hardware and programming errors when they do occur, such that system terminations are avoided more frequently than when another operating system, such as OS MVT, is used.

Specifically, the implementation of one virtual storage for each user in OS/VS2 MVS decreases the chance that one user will inadvertently modify virtual storage of another user or critical portions of the control program. Many integrity features designed to prevent program errors are also basic to the OS/VS2 operating system. Several other features that are unique to OS/VS2 MVS are provided to reduce the impact of errors.

First, functional recovery routines (FRRs) for certain operating system functions are provided. An FRR is a tailored recovery routine for a specific system function. It uses data stored during execution of the system function in an attempt to repair that function when a failure occurs, so that system operation can continue without a re-IPL.

Second, OS/VS2 MVS eliminates many system terminations that would otherwise occur as the result of subchannel errors and control unit lockups by issuing the System/370 CLEAR I/O instruction to reset byte and block multiplexer subchannels when necessary.

Third, OS/VS2 MVS supports operator-initiated recovery via the restart function on the operator frame of the 3036 Console. If the operator suspects the system is in a loop or uncoded wait state, instead of re-IPLing and restarting the system, he can select the restart function, which gives control to the recovery termination manager in OS/VS2 MVS. The recovery termination manager initiates normal recovery processing, using FRRs and other recovery routines to recover the operating system without a system termination.

OS/VS2 MVS multiprocessing support also includes additional recovery facilities designed to further improve the availability offered by a 3033 tightly coupled multiprocessing configuration. When a hard processor failure occurs in one processor, the alternate processor recovery (ACR) facility is invoked in the functional processor to recover from the processor failure so that system operation continues uninterrupted without the failing processor. Channel set switching hardware in the 3033 processors is used to recover I/O operations in progress in the system with the failing processor.

After processor and I/O recovery have been performed in a tightly coupled multiprocessing environment, OS/VS2 MVS automatically varies offline the malfunctioning processor and all channel paths to it. If the primary system console was attached to the failing processor, an automatic switch to the console of the functioning processor is made by OS/VS2 MVS.

Less Complex Operational Requirements

A tightly coupled multiprocessing configuration has less complex operational requirements than two uncoupled systems because it presents a single system image to the operator even though there are two systems in the configuration. The operator has one operational interface to the entire system, one job scheduling interface, and one point of control for all the resources in the configuration. Also, the operator must communicate with and control only one control program instead of two.

Improved Resource Utilization

Resource utilization in a tightly coupled multiprocessing configuration is improved over that of two uncoupled systems because load leveling occurs between the two systems, there is a reduction in the amount of processor storage required by the resident control program, there is I/O device pooling, and the need for using Shared DASD support is eliminated.

Load leveling occurs for processor processing and I/O processing because of the way in which OS/VS2 MVS can schedule task execution and I/O operations in a tightly coupled multiprocessing configuration. Load leveling reduces the peak and valley periods of processor and I/O utilization that normally occur in two uncoupled systems, as follows.

The two processors are considered to be system resources that, when available, are allocated to ready tasks. Usually, either processor is capable of processing each task in the system. Thus, as soon as a processor becomes available, it is allocated to the highest priority queued ready task. Since there are on the average twice as many tasks in a multiprocessing configuration as in one system in a two uniprocessor environment, the chances that no task in the multiprocessing system is ready to execute and available processor time will be unutilized are significantly reduced. An I/O operation is begun on any available channel path in the total configuration and is not limited to being started only on the channel path(s) in one processor.

Since there is only one copy of the OS/VS2 MVS multiprocessing control program resident in processor storage in a tightly coupled multiprocessing configuration, more processor storage is available for paging (which can benefit performance) than in two uncoupled systems with the same total amount of processor storage as in the multiprocessing configuration, each of which has an OS/VS2 MVS uniprocessor control program resident.

I/O devices with programmable channel- or string-switching features installed are pooled in a tightly coupled multiprocessing configuration for use by both processors. More than half the total number of switched I/O devices can be allocated to an individual job when necessary. The pooling of I/O devices and sharing of processor storage permits the execution of jobs with larger processor storage and I/O device requirements than is possible using one system in a configuration with two uncoupled systems.

The sharing of I/O devices and processor storage also enables the OS/VS2 MVS control program to automatically handle peak load situations within jobs and to balance the workload across systems. Manual balancing of the workload between two systems, as is required for two uncoupled systems, is not required for a tightly coupled configuration.

The pooling of I/O devices can reduce the total number of I/O devices required in a tightly coupled multiprocessing configuration when compared with the I/O device requirements for two uncoupled systems that are to handle the same large I/O job or peak load direct access storage requirements.

Since there is only one control program, there is no need to split any data base into two parts, one for each system, or to use Shared DASD support in order to share a data base between the two systems, as is required for two uncoupled systems. The use of Shared DASD support results in reduced throughput for two uncoupled systems because of the interference it introduces. This throughput reduction is not incurred in a tightly coupled multiprocessing configuration since there is only one control program and it can maintain the integrity of a shared data base without using Shared DASD support.

Operational Flexibility

A tightly coupled 3033 multiprocessing configuration can be divided into two systems that operate in uniprocessor mode when this is required to handle special environment situations. For example, a uniprocessor mode system might be required (1) for planned preventive maintenance, (2) as a test system for a system programmer, or (3) to run a programming system other than OS/VS2 MVS (such as VM/370).

The two systems in a multiprocessing configuration can be divided such that only the hardware components (processor storage and I/O devices) actually required for the special environment system are allocated to one processor, leaving the balance of the hardware

resources available for normal production processing. When two uncoupled systems are present in an installation, one total system (half the total hardware resources) must be allocated to the special environment system regardless of its actual processor storage and I/O device requirements (except for any switched I/O devices).

In addition to being able to tailor the resources of a special environment uniprocessor mode system, an installation can perform the reconfiguration dynamically in a multiprocessing configuration without the necessity of stopping the system, canceling or interrupting jobs, or quiescing the system.

In an installation with two uncoupled systems, production operations in one system must be stopped or allowed to quiesce before the system can be used for the special environment operation. Then production processing must be restarted after the special processing is completed. Thus, productive capability is lost during the time the system is being quiesced and later during restart operations. This productive capability is not lost in a tightly coupled multiprocessing configuration.

Note also that there is very little possibility of losing half of the hardware components of a tightly coupled multiprocessing configuration as a result of one failure. If a processor in one system fails, processor storage and switchable devices in that system can still be used by the other processor. If all processor storage in one system fails, the processor, channels, and I/O devices in that system are still available to the multiprocessing configuration. In an environment with two uncoupled systems, a failure in one system causes a loss of the entire system except for any I/O devices that can be switched to the other system.

Improved Growth Options

The installation of tightly coupled multiprocessing in an environment with two uncoupled 3033 systems to handle added workload or the addition of an application is an alternative, less expensive growth step than the installation of a third uncoupled system.

Assume the two uncoupled systems can handle the current workload but the installation of another data base application or more terminals, for example, would cause unacceptable performance during peak load periods. Assume also that peak load periods in the two systems occur at different times and one or both processors do not have very high processor utilization during non-peak periods.

Installation of tightly coupled multiprocessing would enable the resources of both systems to be utilized more efficiently during peak load periods and could provide acceptable performance because of the load leveling that occurs. If the workload continued to grow and additional resources were required, a third system (Model 145, 3031 Processor, or higher) could be installed and loosely coupled with the existing tightly coupled configuration to provide the desired performance.

Improved Throughput Possibilities

As a result of the sharing of common resources that occurs in a tightly coupled multiprocessing environment, its internal performance is generally less than twice that of a single processor environment with similar resources. Thus, the internal performance of a 3033 Multiprocessor Complex is generally in the area of 1.6 to 1.8 times that of a 3033 Processor Complex running with MVS and identical programs. However, because of the load leveling that occurs in a tightly coupled

configuration, the throughput achieved for a given job stream can be greater than that achieved when the same job stream is split and processed by two uncoupled systems with total resources equal to those of the two coupled systems.

The reason the throughput potential is greater for the tightly coupled configuration is that load leveling enables most processor and I/O time that is unoverlapped when two job streams execute in two uncoupled systems to be overlapped when the same two job streams execute in a tightly coupled configuration, as follows.

In a single system environment, unoverlapped I/O time occurs when there are no tasks ready to execute until some I/O completes and unoverlapped processor time occurs when no I/O is operating and none can be started until certain instruction processing completes. When two job streams execute in two uncoupled processors, unoverlapped processor and I/O time occur in both systems and the chances that unoverlapped processor time or I/O time occurs simultaneously in the two systems are small.

When these two job streams execute in a tightly coupled multiprocessing system, because of the way work is dispatched (in effect, each processor can process the other processor's job stream), most unoverlapped processor time that occurs in one system in the uncoupled environment is overlapped with unoverlapped I/O time that occurs in the other system in an uncoupled environment and vice versa. Studies and installation experience have shown that when the two job streams involved are relatively unbalanced as far as processor utilization is concerned (80 percent for one and 50 percent for the other, as an example, instead of both 80 percent), the potential for increased throughput in a tightly coupled configuration, as a result of load leveling, is greater.

Greater throughput potential for a tightly coupled multiprocessing configuration also results from the other advantages such a configuration offers, such as the (1) availability of more processor storage for paging (because one multiprocessing MVS control program requires less processor storage than two uniprocessor MVS control programs), (2) elimination of Shared DASD support, (3) elimination of lost productivity as a result of quiesce and restart operations, (4) availability of more of the total system resources for production processing when a special uniprocessor mode system is required, and (5) reduction in job reruns, data set rebuilding, and re-IPLing as a result of the improved availability features that reduce the number of system terminations.

The performance of an individual application in a tightly coupled multiprocessing environment versus that achieved in a uniprocessor system is related to the amount of multitasking within the application and how well the processing is balanced among tasks. Best performance is achieved when the work required to process a given transaction is distributed fairly evenly across the tasks that are to process the transaction. Good performance can be achieved when no task performs more than 50 percent of processing for a transaction.

ADVANTAGES OF LOOSELY COUPLED MULTIPROCESSING CONFIGURATIONS

Like tightly coupled multiprocessing configurations, loosely coupled multiprocessing configurations offer advantages over single system and uncoupled multiple system configurations, such as higher availability (via better recovery techniques and concurrent maintenance) and operational efficiency and flexibility (via a single input stream, a single operator interface, pooled I/O resources, automatic workload balancing, and better handling of peak load job conditions).

In addition, however, loosely coupled multiprocessing configurations offer certain advantages over tightly coupled multiprocessing configurations, as follows:

- Larger growth potential and growth without disruption. Many more systems can be coupled in a loosely coupled than in a tightly coupled configuration and additional systems can be added with relatively little disruption to operations. Because System/370 tightly coupled multiprocessing is limited to a maximum of two systems, the addition of a third system can be accomplished only by going to a loosely coupled configuration.
- Loosely coupled systems can be asymmetric. Different System/370 processors with unlike hardware characteristics (processor speed, number of channels, and processor storage size) can be loosely coupled. Although processor storage symmetry is not required for 3033, Model 168, and Model 158 Model 3 tightly coupled multiprocessing configurations, processor symmetry (two 3033, 3158, or 3168 processors) is required. For a Model 158 Model 1, processor storage size symmetry is also required unless an RPQ that permits asymmetric processor storage configurations is installed.
- A larger range of models can be loosely coupled. JES3 will support 3033, 3032, and 3031 uniprocessor configurations; 3033 and 3031 attached processor configurations; 3033 multiprocessor configurations; Model 145, 155 II, 158, 165 II, and 168 uniprocessor configurations; and Model 158 Multiprocessing, 158 Attached Processor, 168 Multiprocessing, and 168 Attached Processor Systems (running under OS/VS2 MVS) in a loosely coupled configuration. The configuration can range from a minimum of one Model 145 with 1024K of processor storage to a maximum of eight 3033 tightly coupled multiprocessing systems that are loosely coupled.

A tightly coupled multiprocessing configuration has one significant advantage over a loosely coupled configuration. While either system can handle peaks in the number of jobs processed, only a tightly coupled system can handle peaks that occur within a single multitasking job (such as an online data base application), since the processor, processor storage, and I/O devices of both systems can be used for the job during peak requirements. The unique advantages of both types of multiprocessing configurations can be obtained when required by combining the two within an installation.

In addition to the general benefits of tightly and loosely coupled multiprocessing, 3033 tightly and loosely coupled multiprocessing configurations offer large-system users the following specific advantages:

- Compatible growth for Model 65, 155, 155 II, 158, 165 II, and 168 uniprocessor systems and Model 158 and 168 Attached Processor and Multiprocessing Systems
- Compatible growth for multiple-system installations with System/360 and/or System/370 processors
- A growth path for Model 165 II and purchased Model 165 and 168 uniprocessor systems
- All the new availability and serviceability features provided in a 3033 uniprocessor configuration, which are of even greater benefit in a multiprocessing environment
- Enhancements to the channel recovery hardware provided in Model 168 tightly coupled multiprocessing systems and operational convenience provided by the elimination of the configuration control panel of the Model 168

- Support of all the hardware and programming systems features provided for 3033 uniprocessor systems operating in EC and DAT modes, including support of the System/370 Extended Facility during multiprocessor mode operations

30:10 3033 MULTIPROCESSING ARCHITECTURE AND HARDWARE

3033 PROCESSOR COMPLEX AND MULTIPROCESSOR COMPLEX HARDWARE DIFFERENCES

The following identifies the major areas of architectural and hardware implementation differences between a 3033 Multiprocessor Complex and a 3033 Processor Complex (that is, facilities implemented in tightly coupled multiprocessing but not uniprocessor configurations):

- 3038 Multiprocessor Communication Unit - a required unit (discussed in Section 30:05) that provides the communication between the two processors and consoles in a 3033 Multiprocessor Complex.
- Prefixing - a method of assigning unique areas of processor storage to addresses 0 to 4095 for each processor
- Processor addressing and STORE CPU ADDRESS instruction - required to specifically identify each processor
- Time-of-day clock - synchronization of the two physical clocks to provide one logical clock for the multiprocessing configuration
- Interprocessor programmed communication (SIGNAL PROCESSOR instruction) - required to enable a processor to request services of the other processor and to alert it to conditions to which it must respond during multiprocessing mode operations. For example, this capability is used during the initialization of multiprocessing mode operations, for reconfiguring hardware components, and in recovery procedures that occur after a processor failure.
- Interprocessor hardware communication - required to alert a processor to conditions in the other processor and to synchronize certain operations in both processors during multiprocessing mode operations
- Channel set switching hardware - used to expand the recovery provided by the OS/VS2 multiprocessing control program after a processor failure occurs during multiprocessing mode operations by enabling the operational processor to control the channels of the failing processor. This function is operative only for multiprocessing mode operations.

The following instructions are included in the multiprocessing feature and are valid in uniprocessor as well as multiprocessing mode:

```
SET PREFIX
STORE PREFIX
STORE CPU ADDRESS
SIGNAL PROCESSOR
```

The processor-dependent machine check extended logout area for the processors in a tightly coupled 3033 multiprocessing configuration contains additional information, such as the configuration in effect and the value in the prefix register at the time of the interruption.

Note that most of the balance of this subsection also applies to the multiprocessing functions in the 3033 Model A Processor and 3042 Attached Processor, which are part of a 3033 Attached Processor Complex.

The text includes references to attached processor mode operations where appropriate.

PREFIXING

Each processor in a tightly coupled multiprocessing configuration must have a unique area of storage to be used for permanently assigned locations and logout areas, which in the 3033 Processor (and 3042 Attached Processor) are contained in addresses 0 to 1928. Since there is only one set of storage locations with these addresses in shared processor storage, a means of assigning these addresses to two different storage areas, one for each processor, is required. Prefixing, sometimes referred to as direct address relocation (not to be confused with dynamic address translation), is the technique used. Prefixing is applied to real storage addresses only, that is, to translated addresses when dynamic address translation is operative.

The implementation of prefixing in a 3033 tightly coupled multiprocessing configuration permits the assignment of addresses 0 to 4095 to any storage area of 4096 bytes that begins at an address that is a multiple of 4096. (Addresses up to 4095 are prefixed so that the prefixed area for each processor can contain certain control blocks that are required for each processor, which, in turn, simplifies control program coding). A 4K-byte storage area that is assigned to contain addresses 0 to 4095 for a given processor is called a prefixed storage area (PSA).

Storage addresses are prefixed by means of a prefix value register. Each processor has its own prefix value register and, for proper system operation, a unique 12-bit prefix value must be assigned for each processor during multiprocessor mode operations. Prefix registers are initialized to zero during system reset and clear.

The SET PREFIX and STORE PREFIX privileged instructions are provided to access the prefix register. SET PREFIX is not retryable. STORE PREFIX is totally retryable. A given processor can set or store only its own prefix register. A processor cannot address the prefix register of the other processor.

SET PREFIX is used to place a 12-bit prefix value in the prefix register in bit positions 8 to 19. Bits 0 to 7 and 20 to 31 of the prefix register are ignored. The translation lookaside buffer is purged when a SET PREFIX instruction is issued. STORE PREFIX can be used to place the prefix value contained in a prefix register in processor storage. Bits 0 to 7 and 20 to 31 are stored as zeros.

Prefixing operates as follows. When a processor references a real storage address in the range of 0 to 4095 (the high-order 12 bits, 8 to 19, of the real 24-bit storage address are zeros), bits 8 to 19 of the prefix register for that processor are added to bits 8 to 19 of the real address. The new address will then point to a location in the PSA of the processor. This is called forward prefixing. When a processor references a real storage address in the 4096-byte block that is pointed to by its prefix register (that is, an address in its own PSA), zeros are substituted for bits 8 to 19 of the real storage address so that an address range of 0 to 4095 results. This is called reverse prefixing. Reverse prefixing essentially provides the capability for either processor to access locations 0 to 4095.

Prefixing is always active and is not subject to mode control. That is, prefixing always occurs in a 3033 Processor that has the multiprocessing feature, whether the processor is in uniprocessor or multiprocessor (or attached processor) mode, and in a 3042 Attached Processor. Prefixing is applied to all real storage references made by

a processor. It is implemented in hardware and does not affect instruction execution time. Channel hardware references to locations 0 to 4095 (for a channel address word, during channel status word storing, etc.) are also prefixed.

References by a channel to a channel command word (CCW), I/O data, indirect data address lists, and machine check extended logout area locations, either during system operation or IPL, are not prefixed. This approach permits channel programs (CCW lists) and I/O buffers that are contained in the PSA of one processor to be executed by the other processor without the necessity of moving the channel program and buffers to the PSA of the other processor. Prefixing is also not performed during a store status operation.

PROCESSOR ADDRESSING

In a 3033 tightly coupled multiprocessing configuration, each processor has a unique four-bit address, which is stored during certain external interruptions to identify the processor involved. The processor address is also used in the SIGNAL PROCESSOR instruction, which is discussed under "Interprocessor Programmed Communication".

Addresses 0000 and 0001 are used for the two processors in a 3033 Multiprocessor Complex (and 3033 Attached Processor Complex). These addresses are established by the customer engineer during system installation (the installation can decide which address is assigned to each processor) and are effective during uniprocessor and multiprocessor (and attached processor) mode operations. The STORE CPU ADDRESS privileged instruction is provided only in multiprocessor (and attached processor) models of the 3033 (and the 3042 Attached Processor) to enable a program to obtain the address of the processor in which it is executing.

An example of a situation in which this instruction is issued is after the IPL of a multiprocessing configuration. IPL from either processor is possible. Therefore, STORE CPU ADDRESS is issued by the control program to determine the processor in which it is operating so that the other processor can be started via execution of a SIGNAL PROCESSOR instruction, which requires a processor address.

STORE CPU ADDRESS should not be confused with STORE CPU ID, which is a System/370 instruction implemented in all 3033 Processors. The STORE CPU ADDRESS instruction is totally retryable.

TIME-OF-DAY CLOCK

Because there are two time-of-day clocks in a tightly coupled multiprocessing configuration, one in each processor, modifications to uniprocessor time-of-day clock hardware are required to ensure that both clocks contain the same time during multiprocessor (or attached processor) mode operations with both processors online. Logically, there must be only one clock in the multiprocessing system. When the two processors in a 3033 Multiprocessor Complex operate as independent processors in uniprocessor mode (or the 3033 Processor in a 3033 Attached Processor Complex operates in uniprocessor mode), these changes are disabled so that the two time-of-day clocks operate independently from one another.

When the two processors are operating in multiprocessor (or attached processor) mode, a SET CLOCK instruction can be executed by either processor, if the time-of-day clock security switch for either processor is held in the enable set position. An enable set condition on one processor is broadcast to the other processor to accomplish this clock-

setting function. When a time-of-day value is set in one clock, the OS/VS2 multiprocessing control program ensures that the identical time is contained in the other clock.

When multiprocessor (or attached processor) mode is in effect, stepping pulses for both time-of-day clocks are provided by only one of the two time-of-day clock oscillators so that both clocks step synchronously.

The values of the two clocks are checked for synchronization by a combination of OS/VS2 multiprocessing support and hardware. OS/VS2 compares bits 0 to 31 of two time-of-day clocks to determine an out-of-synchronization condition in the high-order portion of the clocks. The synchronization of bits 32 to 51 in each clock is checked by hardware by the use of a one-second synchronization pulse that is broadcast by each processor to the other processor.

A time-of-day clock synchronization check external interruption condition is generated (normally simultaneously in both processors) when bits 32 to 51 of the two clocks are out of synchronization. This interruption is maskable by the external mask in the current PSW and bit 19 in control register 0. The interruption condition can be handled by whichever processor enables itself for synchronization checks first.

A synchronization check interruption condition continues to be presented until (1) bits 32 to 51 in the two clocks are equal or differ by less than a certain number of nanoseconds or (2) the latch indicating multiprocessor (attached processor) mode is in effect turns off. This interruption is enabled only during the times that the OS/VS2 control program is checking clock synchronization, such as after IPL, setting the time-of-day clock, and varying a processor online.

So that unique time-of-day values are provided in the event that the processors simultaneously issue a STORE CLOCK instruction, the clock value stored contains zeros in bits 52 and 53 for processor 0 and ones in bits 52 and 53 for processor 1. If a STORE CLOCK instruction is issued for a time-of-day clock that is in the stopped state, bits 52 and 53 are not stored.

Modification of the hardware implementation of the interval timer for operation in a tightly coupled multiprocessing environment is not required. Only one interval timer is used. Unlike the time-of-day clock, the interval timer is located in processor storage (in a PSA) and, therefore, the timer can be addressed by both processors in a 3033 Multiprocessor Complex no matter which one of the two interval timers is used by the multiprocessing control program.

INTERPROCESSOR PROGRAMMED COMMUNICATION

The SIGNAL PROCESSOR (SIGP) privileged instruction is provided in 3033 tightly coupled multiprocessing configurations to support program-initiated communication between the two processors. Its operands indicate the address of the processor being signaled, the function to be performed by the addressed processor (via a one-byte order code), and a register in the signaling processor in which status information from the addressed processor can be placed, if necessary. A processor can signal itself by placing its own processor address in the SIGP instruction when uniprocessor or multiprocessor (or attached processor) mode is in effect.

When a SIGP instruction is executed, the addressed processor receives the signal via the hardware or 3036 Console interface between the two processors. The addressed processor must be operating in multiprocessor (or attached processor) mode in order to receive a SIGP instruction.

The addressed processor decodes the order to be performed, performs the operation if possible, and sends a response back to the signaling processor to indicate the action taken. The response consists of a condition code setting and, in some cases, a set of status bits.

Status bits are presented in response to a SIGP instruction when the addressed processor cannot perform the indicated function (because another external call is pending, the addressed processor is in the stopped or check stopped state, the operator is performing an alter or display operation, for example). The SIGP instruction is not retryable.

The orders that can be specified in the SIGP instruction are the following:

Sense. The addressed processor is requested to respond with an indication of its status (operating, stopped, not operational, not ready, external call pending, resetting, for example). A sense could be issued by a processor at IPL, for example, to determine whether the other processor is online.

External Call. An external call type of external interruption condition is generated in the addressed processor, and is taken if the addressed processor is enabled for these interruptions (via bit 18 in control register 0 and current PSW bit 7). If the addressed processor is disabled for this interruption, the external call interruption remains pending. Only one such interruption for a given processor can be pending at a time.

External call is issued by a processor to request that a service be performed by the other processor. The act of one processor in a multiprocessing configuration sending a request to the other processor via programmed communication is called "shoulder tapping".

The external call and emergency signal orders cause the processor address of the processor that issued the SIGP instruction to be stored in processor storage locations 132 and 133 when the external interruption occurs in the addressed processor.

Emergency Signal. An emergency signal external interruption condition is generated in the addressed processor and is taken if the addressed processor is enabled for these interruptions (via bit 17 in control register 0 and current PSW bit 7). The interruption remains pending if such interruptions are disabled. Only one emergency signal interruption can be pending at a time.

Emergency signal, like external call, is issued by one processor when a request is to be communicated to the other processor. Emergency signal and external call functions are used to differentiate between two categories of requests. Emergency signal is used in OS/VS2 multiprocessing support primarily to signal a processor failure via programming and to coordinate purging of the TLBs.

Start. The addressed processor performs the same function as when its start key is pressed. The processor whose 3036 Console was used for the IPL issues a start to the other processor, for example, during initialization of the configuration for multiprocessor (or attached processor) mode operations.

Stop. The addressed processor performs the same function as when its stop key is pressed. A processor issues a stop to the other processor, for example, when the addressed processor has been varied offline by the operator.

Restart. The addressed processor performs the same function as when the PSW reset function is invoked via the operator frame. The restart occurs without a reset and the prefix register is not reset.

Initial Program Reset. The addressed processor, its dedicated channels, and attached dedicated and shared I/O devices are reset (as per a system reset without a system clear). The control registers are initialized in the addressed processor, and its PSW register, prefix register, CPU timer, and clock comparator are set to zero. The reset signal is not propagated to the other processor.

Program Reset. The addressed processor, its dedicated channels, and attached dedicated and shared I/O devices are reset (as per a system reset without a system clear). The reset signal is not propagated to the other processor. This order can be used during recovery operations to reset a hung I/O condition.

Stop and Store Status. The addressed processor enters the stopped state and the store status function is invoked. This order can be used to enable the functional processor to attempt to obtain from the failing processor the CPU timer, clock comparator, PSW, prefix, general register, and floating point register values, which are needed for a successful recovery.

Initial CPU Reset. An initial program reset, except for channel resetting, is performed in the addressed processor. Pending channel interruptions are not reset. This function can be performed only via the SIGP instruction.

CPU Reset. A program reset, except for channel resetting, is performed in the addressed processor. Pending channel interruptions are not reset. This function can be performed only via the SIGP instruction.

INTERPROCESSOR HARDWARE COMMUNICATION

Hardware-initiated communication between two processors in multiprocessor (or attached processor) mode occurs as follows:

Malfunction Alert. Whenever a processor loses power or a machine check error causes the processor to enter the check stopped state, a malfunction alert indication is sent to the other processor. A malfunction alert external interruption condition is generated in the receiving processor if multiprocessor (or attached processor) mode is in effect and the interruption occurs if the processor is enabled for this condition (via bit 16 in control register 0 and current PSW bit 7). The address of the malfunctioning processor is stored in processor storage locations 132 and 133.

Time-of-Day Clock Facilities. These functions have already been discussed. Briefly, hardware communication resulting from use of the time-of-day clock consists of (1) broadcasting a time-of-day clock security switch enable set condition in one processor to the other processor, (2) broadcasting the time-of-day clock oscillator pulse from one processor to the other processor, (3) broadcasting a synchronization pulse between the two processors to check clock synchronization in the low-order bit positions (32 to 51), and (4) broadcasting a time-of-day clock synchronization check indication if a synchronization error occurs.

Buffer Intercommunication. The high-speed buffer storage controls in each processor must communicate with each other to ensure that all real storage references by both processors result in access to the most current copy of the addressed data. In addition, changes to the storage

protect keys must be provided to both processors. Therefore, there is a hardware interface between the two buffer controls by which each buffer control broadcasts to the other the real storage addresses of the data changed and the storage protect keys set.

In a two-processor multiprocessing (or attached processor) configuration, buffer contents are managed as follows. When a processor makes a request for data that is not contained in its high-speed buffer, the data is fetched from real storage, placed in that processor's buffer, and sent to the processor, just as in a uniprocessor environment. The high-speed buffer in the other processor is not affected. When a processor, say processor 0, stores data in real storage, its buffer is updated as well if a copy of the data is currently being maintained in the buffer (as in a uniprocessor environment). In addition, however, the real storage address of the data stored is broadcast to processor 1.

Processor 1 then determines whether the contents of that storage location are currently being maintained in its buffer, and if they are, the data in the buffer in processor 1 is marked invalid. This procedure ensures that the next fetch request for the data by processor 1 will cause a real storage fetch and access to the most current version of the data.

Similarly, when a channel in either processor (or the 3033 Processor in an attached processor configuration) stores data in real storage, the address of the data is broadcast to the other processor so that each processor can inspect its own buffer to determine whether the data is currently being maintained there. If so, the data is invalidated.

Both processors can fetch data from their respective buffers simultaneously. During uniprocessor and multiprocessor (or attached processor) mode operations, a processor can fetch eight bytes from its buffer in 114 nanoseconds (two processor cycles). This is called a local request. During multiprocessor (or attached processor) operations, however, if a local request for a buffer fetch occurs simultaneously with a remote request for marking buffer information invalid, the local request is held for 57 nanoseconds (one cycle) while a search for the data is made and another 57 nanoseconds if invalidation must be performed. Hence 171 or 228 nanoseconds (three or four cycles) are required in this case.

Whenever a processor issues a SET STORAGE KEY instruction, its TLB is inspected and any entries that have the same real address as the address specified in the SSK instruction are invalidated, as in a uniprocessor environment. In addition, the absolute address (prefixed real address) is sent to the other processor so that appropriate entries in its TLB can be invalidated.

Interlock Mode for CS, CDS, and TS Instructions. When multiprocessor (or attached processor) mode is in effect, an interlock mechanism for the COMPARE AND SWAP (CS), COMPARE DOUBLE AND SWAP (CDS), and TEST AND SET (TS) instructions is activated. When a processor issues one of these instructions, a portion of the address (bits 23 to 28) of the interlocked location is sent to the other processor so that the other processor cannot store into the location indicated in the instruction between the fetch and store operation the instruction requires. That is, the processor that issued the instruction enters interlock mode.

If one processor, say 1, is in interlock mode and the other processor (0) also enters interlock mode, processor 1 is also prevented from performing any fetch operations until processor 0 leaves interlock mode (completes its fetch and store operation).

30:15 RECORDING AND DIAGNOSTIC PROGRAMS

The same diagnostic programs (ST370, etc.), microdiagnostics (fault locating tests), and recording programs (RMS, SEREP, EREP, OLT, LOA, etc.) are provided for 3033 multiprocessor as for 3033 uniprocessor configurations and are updated as appropriate for multiprocessing hardware. Additional microdiagnostics are provided to test multiprocessing hardware.

When a malfunction occurs in a processor in a 3033 Multiprocessor Complex, the appropriate microdiagnostics provided for a uniprocessor configuration should be executed first. A minimum system consisting of the malfunctioning processor, 2M-bytes of processor storage, one channel group and required I/O device(s), and the 3036 Console is required to execute these microdiagnostics and uniprocessor mode must be in effect. These components can be logically and physically removed from the multiprocessing configuration without quiescing system operations, as discussed previously. The other system can continue to perform production processing while microdiagnostics are being executed in the malfunctioning processor.

If the uniprocessor microdiagnostics do not locate the malfunction, additional diagnostics that test multiprocessing hardware should be executed next.

30:20 PLANNING CONSIDERATIONS

Successful installation and operation of a shared storage multiprocessing installation requires consideration of some factors that need not be considered in a uniprocessor environment or that are of less importance. Additional planning considerations are discussed in this subsection.

PLANNING FOR MAXIMUM SYSTEM AVAILABILITY

System availability is directly affected by the reliability of the components of the system (both hardware and programming) and by the serviceability of the system components. The reliability of a component is defined in terms of its frequency of outage (solid failures in a given time period) while serviceability is measured in terms of diagnostic facilities available and the duration of time required for repair.

The four components of a system that affect its net productive time are the hardware configuration, the environment in which the system operates, the operating personnel, and the operating system (control and processing programs). The following discusses the reliability and serviceability features of these components as they relate to a 3033 shared storage multiprocessing configuration and the steps that can be taken by system designers to maximize the availability of such a system.

Hardware Configuration

The reliability of the hardware components of a 3033 multiprocessing configuration (like that of a uniprocessor configuration) is enhanced by the implementation of extensive hardware retry, programmed retry, automatic hardware deletion, programmed hardware removal, and programmed recovery procedures that prevent system outages because of intermittent or solid failures in many cases.

Hardware serviceability is improved by extensive error recording, of intermittent as well as solid errors, and by the availability of inline

diagnostics, online diagnostics, and enhanced fault-locating microdiagnostics. In addition, as discussed previously in Section 20:35, the 3036 Console provides serviceability features for a 3033 configuration that are not provided for the Model 168.

While system design cannot improve the inherent reliability of any hardware component, a shared storage multiprocessing system can be designed such that the failure of one or more hardware components has minimal impact on the critical subsystem. At the least, critical hardware components should be duplexed. Critical hardware components are those that make up the minimal critical subsystem--that portion of the total hardware configuration that is necessary for the performance of productive work.

For maximum availability, processor features, channels, and I/O devices should be symmetrically configured on the two processors and channel-switching or string-switching features should be installed on control units, where available, to provide access to I/O devices from both processors. Control unit redundancy (cross-channel switching for tape units, 2844 Auxiliary Storage Control for the 2314, string switching for 3330-series, 3340, and 3350 disk storage, etc.) should be used for critical I/O devices to eliminate the loss of access to critical I/O devices because of a malfunctioning control unit.

System designers can plan to use standard and optional serviceability features that are provided, such as OLTEP. When system usage is being planned, adequate time should be allocated for preventive maintenance procedures. In a shared storage multiprocessing environment, maintenance can be performed on one processor while the other processor continues processing (assuming that the system doing productive work consists of the minimal critical subsystem).

Operating Environment

The hardware components of the system configuration have been designed to operate within certain environmental constraints of temperature, humidity, cleanliness, etc., as detailed in System/370 Installation Manual - Physical Planning (GC22-7004). Any unit or system installed in an environment that violates any of these constraints can be expected to experience increased error occurrence and unavailability. Hence system design should ensure that the required physical environment is provided and maintained to avoid outages caused by improper physical surroundings.

Operating Personnel

Knowledgeable operators are highly desirable in any data processing installation. However, in a high-availability environment, having well-trained operators is of even more significance because operator errors or lack of proper responses to the changing operational environment can severely impact system availability.

Although the operator has a single interface to both systems in a 3033 multiprocessing configuration and less action is required for operation of this configuration, as compared to the operational requirements of two independent 3033 uniprocessor configurations, successful operation of a multiprocessing configuration requires a little more knowledge than is required to run one 3033 uniprocessor configuration.

The major new operational procedures that the operator must learn are those associated with hardware reconfiguration, that is, use of the multiprocessing configuration frame of the 3036 Console and the

configurability commands and procedures associated with OS/VS2 multiprocessing support.

An installation-designed operator checklist form for the multiprocessing configuration frame is a highly desirable operator aid. These checklist forms, marked with the most frequently used configurations, should become a standard part of job setup instructions for the operator.

Operating System

The reliability of the control and processing programs that are included in the operating system affects the availability of the system. As discussed previously, higher reliability and availability than provided by previous operating systems were major design objectives of OS/VS2 MVS. Processing program development procedures for the installation should include thorough processing program testing and use of available debugging and problem determination aids as part of the standard application development procedures. These procedures are even more important in the development of routines that interface with the control program.

Despite extensive program testing, however, it is not always possible to test every condition that can arise during the execution of a program. Hence, processing program errors will occur that can cause termination of system operations, particularly during the initial use of any program. Therefore, the use of restart (such as advanced checkpoint restart and warm start) procedures should be preplanned so that system operation can be restarted as quickly as possible after an abnormal termination.

SECTION 40: THE 3033 ATTACHED PROCESSOR COMPLEX

40:05 GENERAL DESCRIPTION

The 3033 Attached Processor Complex is a configuration in which an additional instruction execution capability is provided. A 3033 Attached Processor Complex is a tightly coupled multiprocessing configuration in which two instruction streams can be executed simultaneously. It contains two processors (a Model A 3033 Processor and a 3042 Attached Processor) that share all processor storage in the 3033 Processor and execute under the control of a single multiprocessing operating system. This configuration is supported by OS/VS2 MVS and VM/370.

The 3033 Attached Processor Complex provides a growth path for 3033 uniprocessor users who require additional internal performance, but do not require all the advantages offered by a 3033 Multiprocessor Complex, and offers advantages over two uncoupled 3033 uniprocessor configurations. A 3033 Attached Processor Complex operating under OS/VS2 MVS is capable of providing internal performance 1.6 to 1.8 times that of a 3033 Processor Complex. The internal performance improvement realized when a given 3033 uniprocessor configuration is upgraded to a 3033 attached processor configuration is dependent on the amount of multiprogramming that can be achieved.

The multiprocessing functions implemented in the 3033 Attached Processor Complex are compatible with those implemented in the Model 168 Attached Processor System. Therefore, programs that operate in a Model 168 attached processor configuration can operate in a 3033 attached processor configuration without modification, subject to the constraints listed in Section 10:05. The 3033 Attached Processor Complex is also compatible with the 3031 Attached Processor Complex and provides availability features not offered by Model 158, Model 168, and 3031 attached processor configurations.

COMPONENTS

The components of the 3033 Attached Processor Complex, as shown in Figure 40.05.1, are:

- One 3033 Model A Processor. Functionally, this unit is a uniprocessor model of the 3033 Processor with tightly coupled multiprocessing hardware like that implemented in multiprocessor models of the 3033 Processor.
- One 3042 Attached Processor (AP) physically connected to the 3033 Model A Processor via one 3038 Multiprocessor Communication Unit. The 3042 AP contains an instruction processor function similar in capability to that in the 3033 Model A Processor. The 3042 does not contain any processor storage or channels. The physical size of the 3042 AP, as shown in Figure 40.05.1, is smaller than that of the 3033 Model A Processor, since the 3042 AP does not contain channel and processor storage frames. The 3038 provides a communication path between the two processors and the two 3036 Consoles, as well as processor storage addressing capabilities for the 3033 and 3042.
- Two 3036 Consoles, one for the 3033 Processor and one for the 3042 AP. The inclusion of a 3036 Console for the 3042 AP provides independent power control and power monitoring for the 3042. This

3036 Console can also be used to execute limited diagnostics on a 3042 AP concurrently with operation of the 3033 Processor.

- Two 3037 Power and Coolant Distribution Units, each with a multiprocessing feature installed. Two motor generator sets, one for each processor, are also required.

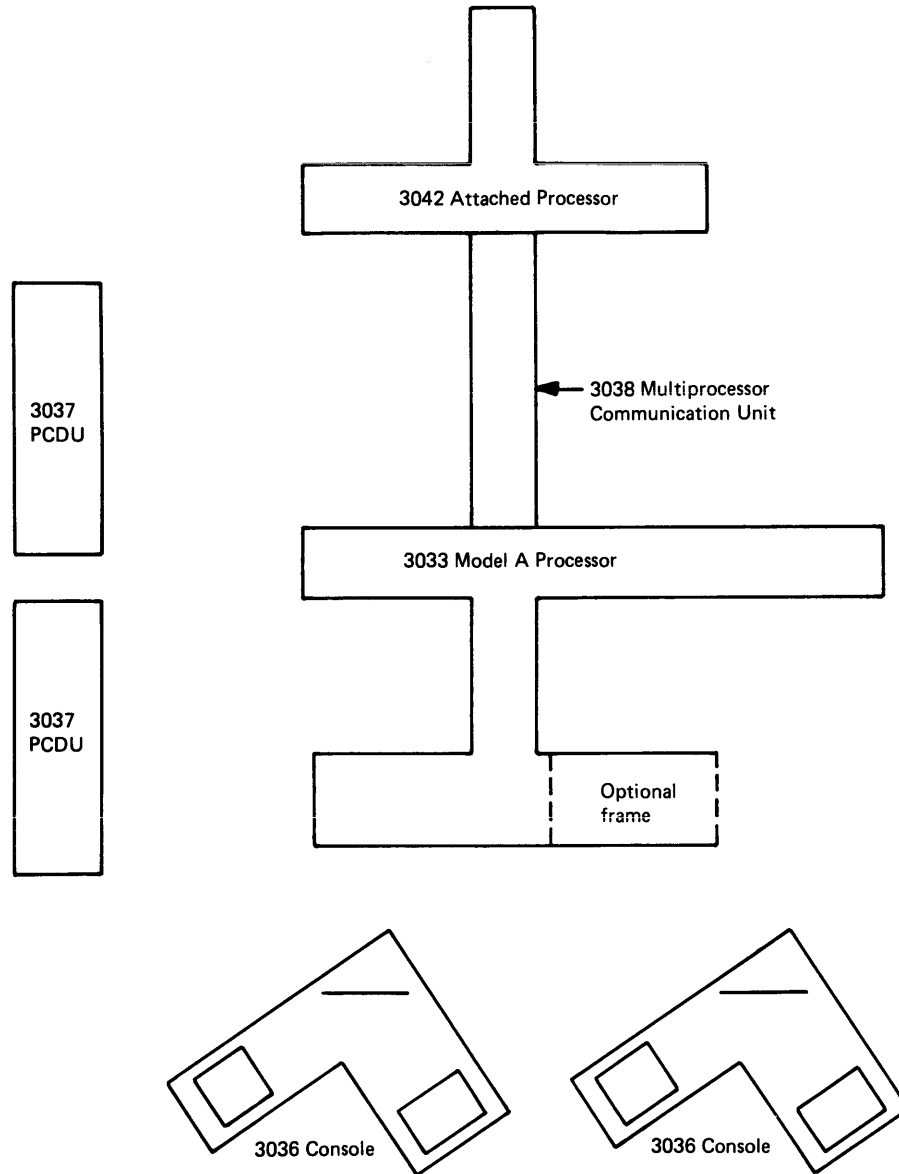


Figure 40.05.1. Components of the 3033 Attached Processor Complex

The 3033 Attached Processor Complex can operate in two modes. When attached processor (AP) mode is in effect, the 3033 Processor and 3042 AP normally operate together as a tightly coupled multiprocessing configuration that shares processor storage in the 3033 Processor. When uniprocessor (UP) mode is in effect, the hardware connection between the 3033 and 3042 is not enabled and only the 3033 Processor can operate as a uniprocessor. Limited diagnostics can be executed on the 3042 when

uniprocessor mode is in effect. Permissible operating configurations are discussed in detail under "3036 Console."

Operating systems other than OS/VS2 MVS and VM/370 that support the 3033 Processor Complex can operate in a 3033 Model A Processor but do not support the 3042 AP (that is, a tightly coupled multiprocessing environment). The performance achieved when these operating systems are used in a 3033 Model A Processor operating in uniprocessor mode will be somewhat less than that obtained using a uniprocessor model of the 3033 Processor.

Note that two 3033 Attached Processor Complex configurations cannot be coupled to form a tightly coupled multiprocessor complex configuration. However, 3033 attached processor configurations can be included in the loosely coupled multiprocessing configurations supported by JES2 Multi-Access Spool and JES3 support in OS/VS2 MVS.

A 3033 Attached Processor Complex can be field converted to a 3033 Multiprocessor Complex. The 3033 Model A Processor must be converted to a 3033 Model M Processor and the 3042 AP must be removed and replaced with a second 3033 Model M Processor.

3033 MODEL A PROCESSOR

The same standard and optional features provided for uniprocessor models of the 3033 Processor are provided for Model A 3033 Processors, with the same limitations and restrictions. However, 3033 Model A Processors also have tightly coupled multiprocessing functions. The A models of the 3033 Processor are basically identical to the M models of the 3033 Processor that are utilized in a 3033 Multiprocessor Complex configuration.

A uniprocessor model (including the no longer available Model U6) of the 3033 Processor can be field-converted to a Model A 3033 Processor and attached to a 3038 Multiprocessor Communication Unit. Once this change is made, the Model A 3033 Processor can still operate as a uniprocessor until the 3042 Attached Processor is installed. When the STORE CPU ID instruction is issued in a 3033 Model A Processor, the processor number 3033 and version code 80 are stored.

Processor Storage

The same eight-way doubleword interleaved processor storage (4M-, 8M-, 12M-, or 16M-bytes) available for 3033 Processors is available for 3033 Model A Processors and is designated as Models A4, A8, A12, and A16, respectively, for the 3033 Model A Processor. Processor storage is divided into eight logical storage elements addressed 0 through 7, each of which can be accessed concurrently.

Nonsimultaneous requests from the 3033 Processor and 3042 AP to a nonbusy logical storage element are handled on a first-come, first-served basis. When both processors simultaneously request access to processor storage, storage priority is granted on a priority-weighted demand basis. That is, a given processor retains access priority to storage and can access storage on successive cycles (rather than every alternate cycle) as long as the processor can sustain storage requests of a higher priority classification than the other processor. Thus, the processor that has storage priority retains this priority until the processor without storage priority has a higher priority request for storage. In this situation, storage priority is switched from one processor to the other.

When tightly coupled multiprocessing functions are present in a 3033 Processor, requests from the processor to its processor storage (local requests) and requests from the 3042 AP to the processor storage contained in the 3033 Processor (remote requests) are channeled to storage via the 3038 Multiprocessor Communication Unit. Therefore, the same amount of time is required for a remote request as for a local request when attached processor mode is in effect and more time is required to access processor storage in a 3033 Attached Processor Complex than in a 3033 Processor Complex.

The processor storage contained in the 3033 Processor in a 3033 Attached Processor Complex is logically divided into 2M-byte units, called storage bands, in order to implement floating storage addressing. Floating storage addressing is implemented so that the physical address range assigned to each band of processor storage can be varied as needed. The address range of each 2M-byte unit of processor storage is specified using a configuration (C2) frame of the 3036 Console (see discussion under "3036 Consoles").

Channel Configuration and I/O Devices

Two channel groups are standard in a 3033 Model A Processor. They are identical to the two standard channel groups in uniprocessor models of the 3033 Processor and the same general channel configuration rules should be observed (see Section 10:25). The same I/O devices can be attached to a Model A as to a uniprocessor model of the 3033 Processor. No channels or I/O devices can be attached to the 3042 AP.

Normally the channels in the 3033 Processor are dedicated to the 3033 and cannot be accessed by the 3042 AP. However, channel set switching hardware is included in the two processors in a 3033 Attached Processor Complex as part of the multiprocessing design. Channel set switching is designed to aid the recovery provided by the OS/VS2 multiprocessing control program after the 3033 Processor fails and must be logically removed from the operational configuration. Channel set switching can be activated by programming or via the 3036 Console (C2 configuration frame).

The Channel Reconfiguration Hardware (CRH) function in OS/VS2 MVS for the 3033 Attached Processor Complex will use channel set switching to permit channel controls in the 3033 Processor to be switched to the 3042 AP under program control without a re-IPL after an uncorrectable 3033 Processor failure occurs. The data paths and processor storage of the 3033 Processor must be functional and power must be on in the units through which the data lines pass; that is, certain 3033 Processor frames, the 3038 unit, and the associated 3036 Console and 3037 Power and Coolant Distribution Unit. Activation of channel set switching permits the 3042 AP to receive and process I/O interruptions from, and initiate I/O operations to, the I/O devices attached to the 3033 Processor. VM/370 attached processor support also utilizes channel set switching, when possible, for recovery after an uncorrectable 3033 Processor failure occurs.

Power Control

Power control for a 3033 Attached Processor Complex is designed to allow the 3042 Attached Processor and certain components within the 3033 Processor to be powered on and off separately. Continued processing may be possible, depending on the component powered down.

Hardware and programming systems support for the 3033 Attached Processor Complex are designed to enable a properly isolated system component to be powered down without the necessity of first quiescing

system operations. This capability enables repair operations to be performed on certain components while normal system operations continue. When the malfunctioning component is repaired, it can be powered on and returned to the functional configuration, again without quiescing system operations.

Each of the components that can be separately powered on and off has its own power control switch that can be set to the local or remote position. The switch must be set to local in order to perform separate power on and power off operations.

The following components within a 3033 attached processor configuration can be powered on and off separately from the other components in the configuration:

- Each director/channel group in the 3033 Processor. The balance of this processor, the I/O devices attached to the powered off channels that are accessible via another channel group, and I/O devices attached to other powered on channels are available for attached processor mode operations.
- The IPPF, execution function, and PSCF of the 3033 Processor. The processor storage of this processor is available but the channels of this processor cannot also be utilized, since the PSCF must be powered on in order to utilize the channel set switching function. Thus, system operation using the 3042 AP is not possible.
- The processor storage in the 3033 Processor. System operation is not possible.
- The half of the 3038 Multiprocessor Communication Unit associated with the 3033 Processor. The 3033 Processor and its associated 3036 Console cannot be used. Thus, system operation is not possible.
- Individual control units and I/O devices attached to the channels of the 3033 Processor. The balance of the system is available for attached processor mode operations.
- The 3042 AP and its associated 3038 half. This enables card changing to be done for the 3042 and its 3038 half while processing continues in the 3033 Processor.

Note that when the 3037 unit for the 3033 or 3042 is powered down, the associated processor, 3036, and 3038 half are also powered down. In addition, when the 3036 is powered down, it also causes a powering down of its associated processor and 3038 half.

For proper system operation, a component (processor, channel group, or I/O device) should be logically varied offline to the operating system before it is powered off. A component must be varied online after it is again powered on if it is to be part of the operational configuration.

When a hardware failure in the 3042 AP prevents its continued operation, the 3042 AP can be varied offline and the configuration can continue operating in attached processor mode with only the 3033 Processor functioning. If uniprocessor mode is established, certain diagnostics can be performed on the 3042 AP using the 3036 Console while normal processing continues in the 3033. At least 2M-bytes of processor storage must be enabled (cross-configured) to the 3042 AP in order to execute the diagnostics. Execution of the complete 3042 AP diagnostic package requires use of the entire 3033 Attached Processor Complex.

The 3033 Model A Processor can operate without power on in the 3042 AP. Thus, the 3042 AP can be powered down separately from the 3033

Model A Processor for card replacement, when necessary. When a malfunction has been repaired in the 3042 AP, it can be powered up and varied online and attached processor mode operations involving the two processors can be resumed. Therefore, when certain types of malfunction occur in the 3042 AP, it may be removed from the functional configuration for diagnostic and repair operations and later returned to the functional configuration without requiring a re-IPL.

When an uncorrectable hardware failure occurs in the 3033 Model A Processor, continued operation in attached processor mode with only the 3042 AP may be possible, depending on the location of the 3033 malfunction. When a malfunction occurs in the instruction processor function of the 3033 Processor, for example, the 3033 can be logically varied offline and channel set switching can be invoked to enable processing to continue without a re-IPL using the 3042 AP. No diagnostics can be executed on the 3033 Processor during processing involving only the 3042 AP.

3042 ATTACHED PROCESSOR

The 3042 Attached Processor, with a 57-nanosecond cycle time, contains an IPPF, execution function, PSCF, and maintenance and retry function, like those of the 3033 Model A Processor. The 3042 AP basically differs from Model A and uniprocessor models of the 3033 Processor in that the 3042 does not contain any processor storage or channels. The PSCFs in the 3042 AP and 3033 Model A Processor communicate with each other to support the sharing of all processor storage available in the 3033 Processor. There are no optional features for the 3042 AP.

The same machine check interruptions are implemented in the 3042 AP as in a 3033 Processor. The length of the processor-dependent machine check extended logout area is the same in the 3042 AP as in a 3033 Processor. The logout data stored for the 3042 is the same as that for the 3033 Model A Processor except that it does not contain anything related to channels or processor storage and does contain additional data in previously unused locations that is unique to the 3042 AP. The processor number 3033 and version code 80 are stored for the 3042 Processor when the STORE CPU ID instruction is issued.

3038 MULTIPROCESSOR COMMUNICATION UNIT

The 3038 Multiprocessor Communication Unit physically connects the two processors in a 3033 Attached Processor Complex. It provides a communication path between the two processors and between the two 3036 Consoles. Functionally, the 3038 is divided in half. Each half is associated with the processor it is attached to and receives its power and water cooling from the 3037 of its associated processor. Each half of the 3038 can be powered up and down separately from the other half and contains an oscillator for timing its associated processor.

The processor storage lines from processor storage in the 3033 Processor and the port 4 (trace function) lines for each 3036 Console are routed through the 3038 unit. Each half of the 3038 also has a set of registers that is used for communication between the two 3036 Consoles. Communication between the two 3036 Consoles is required, for example, to pass status information between the two processors and to execute certain orders of the SIGNAL PROCESSOR instruction.

3036 CONSOLES

The two 3036 Consoles in a 3033 Attached Processor Complex are physically identical to the 3036 Console used in a 3033 Processor Complex. No changes are made to the control panel or the two operating stations. However, the operational diskettes for a 3033 attached processor configuration contain processor and console microcode that is designed to support multiprocessing and interprocessor communication.

The 3036 Console associated with the 3042 AP is cabled connected to the 3042. This cabling provides certain hardwired communication functions between the 3042 AP and its console, such as powering, logging to the diskette, hardware configuring using a configuration display, indicator displays, etc.

Operator to operating system communication utilizing a display station, program access to a diskette drive operating in service record file mode, and remote maintenance operations can also be performed using the 3036 Console associated with the 3042 AP. These operations require the 3036 Console to be attached to a channel in the 3033 Processor, as is the other 3036 Console, and assigned three I/O addresses.

For attached processor mode operations, at least one operating station in each 3036 Console must be functional. This is necessary for power control for each processor and 3036 Console-to-3036 Console communication, which is accomplished via the service support stations. All logouts from a processor are written to the diskette drives of its associated 3036 Console.

In the normal console configuration for attached processor mode operations, one station in each 3036 Console is designated as an operator station while the other is designated as a service support station. One operator station in either 3036 Console is used as the primary operating system console while the other can be defined as a alternate or additional console.

The port assignments for the two stations in the 3036 Console for the 3033 Model A Processor are the same as the port assignments for the two stations in the normal console configuration for a 3033 uniprocessor configuration, as shown in Figure 40.05.2. Port assignments for the two stations in the 3036 Console for the 3042 AP are the same as for the 3033 Model A Processor for components present in both processors.

For the normal configuration, the service support station of the 3036 Console for the 3033 Model A Processor can be utilized to perform the same service operations for its associated processor concurrently with normal processing as in a 3033 uniprocessor configuration, as described in Section 20:30 under "Normal Console Configuration".

The service support station of the 3036 Console for the 3042 AP can be used to perform the following concurrently with attached processor mode processing with the normal console configuration in effect:

- Set up modes of operation for the trace unit in the 3042 AP and enable/disable SSR using the mode control frame
- Display trace event records and associated logouts contained on its system diskette.
- Display processor logouts or system status recording records contained on its system diskette. Any power logouts contained on the mounted diskette can be displayed also.
- Monitor system operation using the operator frame or indicator frames associated with the instruction processor function

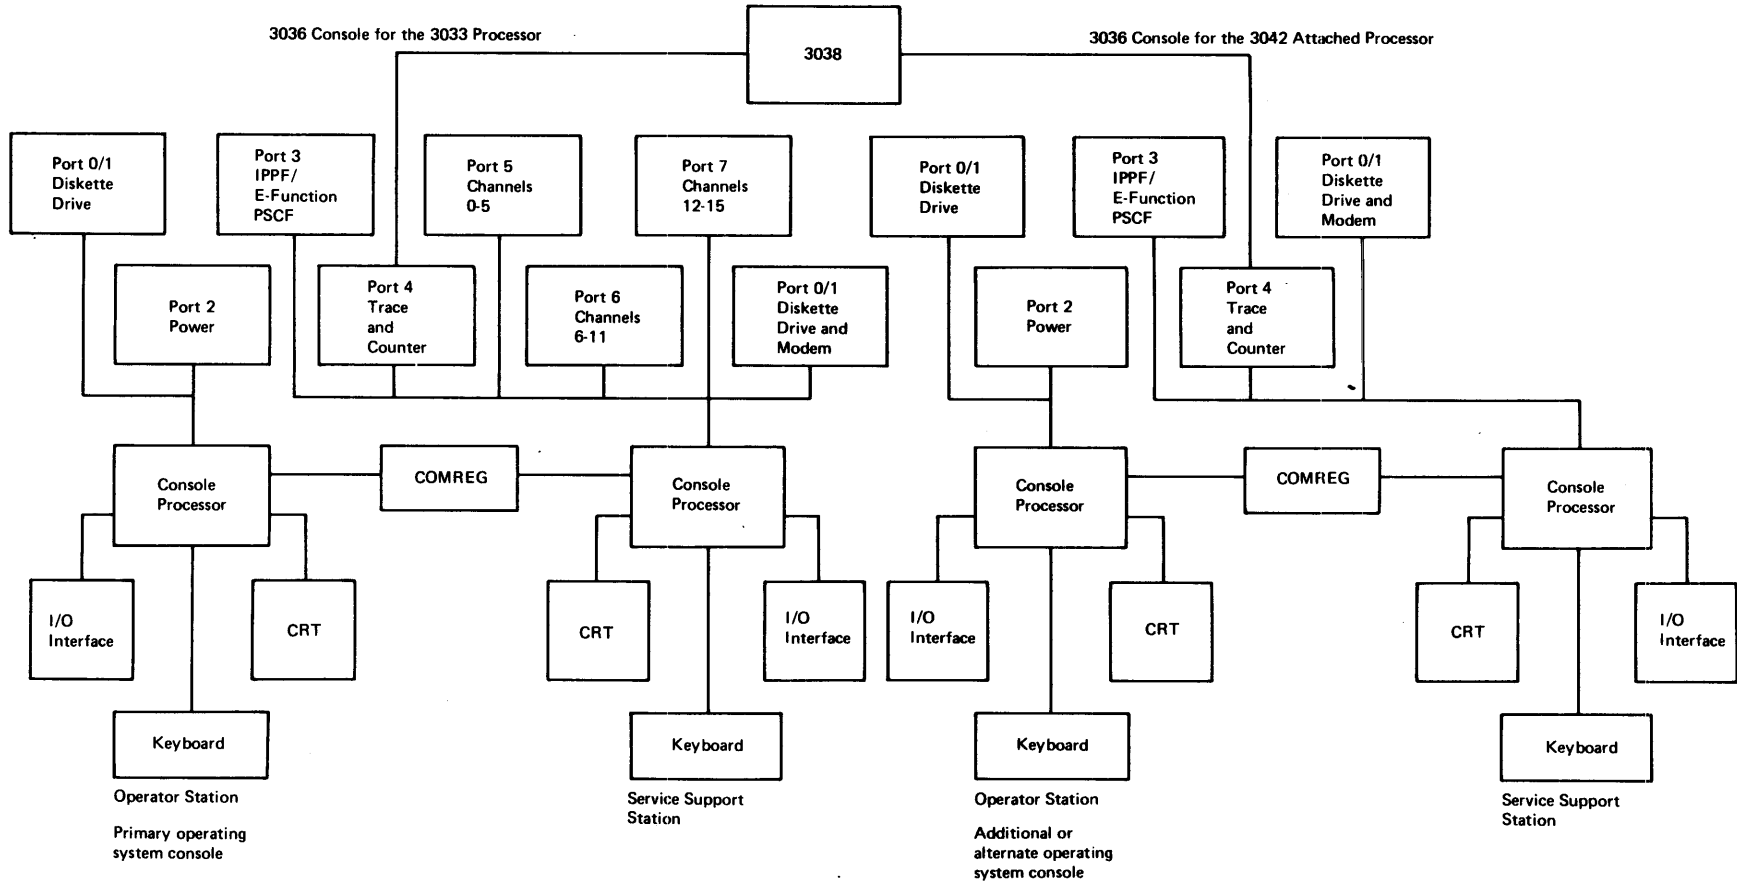


Figure 40.05.2. The normal console configuration for the 3033 Attached Processor Complex

If the operating station being used as the primary operating system console malfunctions, the operator station of the other 3036 Console can be utilized as the primary console. The service support station of the 3036 Console with the malfunctioning station can be assigned as a combined operator/service support station (in <MO> state) and console diagnostics can be run on the malfunctioning console station.

If both operating stations of a 3036 Console are malfunctioning, the associated processor cannot be utilized. If the 3036 Console for the 3042 is the malfunctioning console, processing utilizing only the 3033 Processor is possible. If the 3036 Console for the 3033 Processor is the malfunctioning console, processing using only the 3042 AP is not possible.

The display frames provided for the 3036 Consoles in a 3033 Attached Processor Complex provide the same functions as for a 3033 Processor Complex. Certain frames have been modified and additional frames are provided to support functions required in a tightly coupled multiprocessing configuration. The frames provided for a 3033 Attached Processor Complex are similar to those provided for a 3033 Multiprocessor Complex.

A second configuration (C2) frame, similar to the C2 frame for a 3033 Multiprocessor Complex, is provided for a 3033 Attached Processor Complex to perform configuration functions that in a Model 168 Multiprocessing System are accomplished using the configuration control panel on the 3068 Multisystem Unit. The C2 configuration frame is the only frame that permits an operator to establish configuration information for both processors in a 3033 Attached Processor Complex using either 3036 Console (the one designated as the primary operator console). For other configuring for a processor (such as UCW assignments for the 3033, console designations, powering, etc.), the 3036 Console associated with that processor must be used.

The C2 configuration frame is used to establish the following:

- Uniprocessor or attached processor mode of operation. When attached processor mode is in effect, interprocessor communication hardware is enabled (SIGNAL PROCESSOR interface, high-speed buffer intercommunication, time-of-day clock security switch broadcasting, time-of-day clock synchronization check, malfunction alert signal, and the broadcast of reset functions). These facilities are discussed in Section 40:10.

When attached processor mode is in effect, three operational configurations are possible: the 3033 and 3042 operating together as a tightly coupled configuration (the normal configuration), the 3033 Processor operating alone, and the 3042 AP operating alone.

If the 3042 AP fails, the 3033 Processor can continue operating in attached processor mode until such time as diagnostics are to be performed on the 3042 AP. Then uniprocessor mode must be established. If the 3033 Processor fails in an area not associated with processor storage or the channels, channel set switching can be activated and continued operation in attached processor mode with only the 3042 AP is possible.

When uniprocessor mode is in effect, the interprocessor signals are deactivated. The 3033 Processor can perform normal processing. The 3042 AP is limited to the execution of certain diagnostic routines when uniprocessor mode is in effect.

- Processor storage assignment and interleaving. Each 2M-byte processor storage band in the 3033 Processor has the capability of being enabled for access by both processors, disabled for access by

both processors, or enabled for access by only one of the processors. An address range must be assigned to each enabled band. Eight-way doubleword interleaving or serial operation can also be established, as for a uniprocessor configuration.

For attached processor mode operations, each processor storage band in the 3033 that is enabled to one processor must also be enabled to the other processor. This assignment permits both processors to access all of the enabled processor storage in the 3033 Processor. For uniprocessor mode operations, no processor storage band can be enabled for access by both processors. When uniprocessor mode is in effect and one or more bands of processor storage in the 3033 Processor are enabled for access only by the 3042 AP, processor storage is said to be cross-configured.

Each enabled processor storage band must have a 2M-byte address range assigned. Any of the eight possible 2M-byte ranges (0 to 2M through 14M to 16M) can be assigned to a storage band as long as the same address range is not assigned to more than one enabled band for attached processor mode operations or to two bands allocated to the same processor for uniprocessor mode operations. The address range 0 to 2M must be assigned to one enabled band for attached processor mode operations and to one enabled band for the processor when it is used for uniprocessor mode processing operations; otherwise, a successful IPL cannot be performed.

The address ranges assigned to the storage enabled to the 3033 Processor for uniprocessor mode operations or to shared storage for attached processor mode operations need not be contiguous when an OS/VS2 MVS multiprocessing operating system is used. In addition, address ranges higher than the total amount of processor storage enabled or present in the configuration can be assigned.

For example, address ranges 0 to 6M and 8M to 10M could be assigned to a 3033 attached processor configuration that contained 8M-bytes of processor storage. The addresses between 6M and 8M are marked unavailable by the OS/VS2 multiprocessing control program. The processor storage in the 3033 Processor must be assigned contiguous address ranges for uniprocessor mode operations when an operating system other than OS/VS2 MVS with multiprocessing support is used.

- Oscillator assignments. The system oscillator for each processor in a 3033 Attached Processor Complex is contained in its associated half of the 3038 Multiprocessor Communication Unit and the system oscillator in each processor is disabled. This is done so that the 3033 Processor need not be operational in order for its processor storage to be accessed by the 3042 AP. Each processor still contains its own time-of-day clock oscillator.

For attached processor mode operations, both processors use the same system oscillator and time-of-day clock oscillator so that operations are synchronized. The system oscillator not being used is phase locked to the controlling system oscillator so that control can be switched from the latter to the former without quiescing system operation. For uniprocessor mode, the 3033 Processor can use its own two oscillators or those of the 3042 AP. When processor storage is cross-configured to the 3042 AP to run diagnostics concurrently with processing in the 3033, one set of oscillators must be used by both processors.

- Status of channel set switching. When the normal setting is selected, channel set switching is not active and the 3033 Processor controls the operation of its own channels. When the reverse setting is selected, channel set switching is activated and control of the channels in the 3033 Processor is switched to the 3042 AP.

Note that when channel set switching is activated to switch control of the channels to the 3042 AP, an IPL can be performed utilizing the 3036 Console of the 3042 AP. When channel set switching is not activated, IPL must be performed using the 3036 Console of the 3033 Processor.

- Console(s) to be used for configuration control. The C2 configuration frame can be used to establish control for configuration changes. Changes can be made from both 3036 Consoles using the C2 configuration frame (each 3036 is a master console) or only one 3036 Console (the master). The 3036 Console that cannot make changes (designated the slave console) can only display the configuration in effect using the C2 configuration frame.

When a valid configuration is entered via the C2 configuration frame, each processor is forced to a quiesced state at the end of the execution of its current instruction so that the specified configuration can be made effective. The configuration currently in effect is maintained in storage in the 3036 Console.

When attached processor mode is in effect, the specified configuration becomes effective in both processors simultaneously. When uniprocessor mode is in effect, a change becomes effective in one processor and does not affect the other processor.

As for a 3033 Multiprocessor Complex, processor operations need not be quiesced prior to any alteration of the physical configuration using the C2 configuration frame in a 3033 Attached Processor Complex. Whenever a configuration change is made, all the entries in one or both (as appropriate) high-speed buffers and translation lookaside buffers are invalidated.

The following communication facilities are provided for attached processor mode operations: broadcast of reset functions from one processor to the other, broadcast of time-of-day clock security switch enable set condition, and 3036 Console-to-3036 Console communication.

When a system reset (with or without the system clear) or load (IPL with or without a system clear) is performed for one processor using the operator frame, a corresponding function is also automatically sent (broadcast) to the other processor when attached processor mode is in effect, as shown below.

<u>Function Selected by Operator on Local Processor</u>	<u>Function Performed on Local Processor</u>	<u>Function Broadcast to Remote Processor</u>
System reset (normal)	Program reset	Program reset
System reset (clear)	System clear	Initial program reset
Load (normal)	Program reset	Program reset
Load (clear)	System clear	Initial program reset

The only other communication function provided in the 3036 Console for attached processor mode operations is for the time-of-day clock. When the time-of-day clock security switch for either processor is held in the enable set position, this setting is propagated to the other processor so that the time-of-day clock in each processor is enabled for setting.

40:10 3033 PROCESSOR COMPLEX AND ATTACHED PROCESSOR COMPLEX HARDWARE DIFFERENCES

The 3033 Model A Processor and 3042 AP differ from a uniprocessor model of the 3033 Processor in that they contain tightly coupled multiprocessing functions. This is the only difference between a uniprocessor and Model A 3033 Processor. Differences between the 3042 AP and the Model A 3033 Processor, previously discussed, are also differences between the 3042 AP and a uniprocessor model of the 3033 Processor.

The following identifies the major architectural and hardware implementation differences between 3033 uniprocessor and attached processor configurations (that is, facilities implemented in attached processor but not uniprocessor configurations):

- 3038 Multiprocessor Communication Unit - a required unit (discussed in Section 40:05) that provides the communication between the two processors and consoles in a 3033 Attached Processor Complex.
- Prefixing - a method of assigning unique areas of processor storage to addresses 0 to 4095 for each processor
- Processor addressing and STORE CPU ADDRESS instruction - required to specifically identify each processor
- Time-of-day clock - synchronization of the two physical clocks to provide one logical clock for the attached processor configuration
- Interprocessor programmed communication (SIGNAL PROCESSOR instruction) - required to enable a processor to request services of the other processor and to alert it to conditions to which it must respond during attached processor mode operations. For example, this capability is used during the initialization of attached processor mode operations and in recovery procedures that occur after a processor failure.
- Interprocessor hardware communication - required to alert a processor to conditions in the other processor and to synchronize certain operations in both processors during attached processor mode operations
- Channel set switching hardware - used to expand the recovery provided by the MVS and VM/370 control programs after a 3033 Model A Processor failure occurs by enabling the 3042 AP to control the channels of the 3033 Model A Processor.

The multiprocessing instructions included in the 3033 Model A Processor and 3042 AP, which are valid in uniprocessor as well as attached processor mode, are the following:

SET PREFIX
STORE PREFIX
STORE CPU ADDRESS
SIGNAL PROCESSOR

The SIGNAL PROCESSOR instruction is partially retryable. The other three are totally retryable.

The tightly coupled multiprocessing functions and instructions for the 3033 Model A Processor and 3042 Attached Processor are like the tightly coupled multiprocessing functions and instructions in the 3033 Model M Processors in a 3033 Multiprocessor Complex. See Section 30:10 for a discussion of multiprocessing functions and instructions.

Table 40.10.1 contains a comparison of the hardware features of the Model U 3033 Processor, Model A 3033 Processor, and 3042 AP.

Table 40.10.1. Comparison of hardware features in the Model U 3033 Processor, Model A 3033 Processor, and 3042 AP

Feature	Model U 3033 Processor	Model A 3033 Processor	3042 AP
BC and EC mode of operation	Std	Std	Std
Byte-oriented operands	Std	Std	Std
Channels	Twelve std Four opt	Twelve std Four opt	No
Channel indirect data addressing	Std	Std	No
Channel retry data	Std	Std	No
Channel set switching	No	Std	Std
Channel-to-Channel Adapter	Opt	Opt	No
Clock comparator and CPU timer	Std	Std	Std
Direct Control	Std	Std	Std
Dynamic address translation	Std	Std	Std
ECC on processor storage	Std	Std	No
Expanded machine check interruptions	Std	Std	Std
Extended precision floating point	Std	Std	Std
High-speed buffer	Std-64K	Std-64K	Std-64K
Instruction retry	Std	Std	Std
Interprocessor hardware communication	No	Std	Std
Interprocessor programmed communication (SIGP instruction)	No	Std	Std
Interval timer	Std	Std	Std
Monitoring feature	Std	Std	Std
Prefixing	No	Std	Std
Processor storage	4M-, 8M-, 12M-, 16M-bytes	4M-, 8M-, 12M-, 16M-bytes	No
Program event recording	Std	Std	Std
Reloadable control storage	Std (for instruction processor function and channels)	Std (for instruction processor function and channels)	Std (for instruction processor function only)
Reference and change recording	Std	Std	Std
SSM instruction interruption	Std	Std	Std
Store and fetch protection	Std	Std	Std
Store status function	Std	Std	Std
System/370 Extended Facility	Std	Std	Std
Time-of-day clock	Std	Std	Std
Translation lookaside buffer	Std	Std	Std

40:15 ADVANTAGES OF A 3033 ATTACHED PROCESSOR CONFIGURATION

In addition to increased internal performance over a uniprocessor configuration, an attached processor configuration offers advantages over two uncoupled uniprocessor configurations with the same total resources as the attached processor configuration.

LESS COMPLEX OPERATIONAL REQUIREMENTS

An attached processor configuration has less complex operational requirements than two uncoupled systems because it presents a single system image to the operator even though there are two instruction processor functions in the configuration. The operator has one operational interface to the entire system, one job scheduling interface, and one point of control for all the resources in the configuration. In addition, the operator must communicate with and control only one control program instead of two.

IMPROVED RESOURCE UTILIZATION

Resource utilization in an attached processor configuration is improved over that of two uncoupled systems because load leveling occurs between the two systems, there is a reduction in the amount of processor storage required by the resident control program, all I/O devices in the configuration can be accessed by the 3033 Processor, and the need for using Shared DASD support is eliminated.

Load leveling occurs for the two processors because of the way in which OS/VS2 MVS can schedule task execution in a tightly coupled configuration. Load leveling reduces the peak and valley periods of processor utilization that normally occur in two uncoupled systems, as follows.

The two processors are considered to be system resources that, when available, are allocated to ready tasks. Usually, either processor is capable of processing each task in the system. Thus, as soon as a processor becomes available, it is allocated to the highest priority queued ready task. Since there are on the average twice as many tasks in an attached processor configuration as in one system in a two-uniprocessor environment, the chances are significantly reduced that no task in the attached processor configuration will be ready to execute and available processor time will be unutilized.

Since there is only one copy of the OS/VS2 MVS multiprocessing control program resident in processor storage in an attached processor configuration, more processor storage is available for paging (which can benefit performance) than in two uncoupled systems with the same total amount of processor storage as in the Model A 3033 Processor, each of which has an OS/VS2 MVS uniprocessor control program resident.

While the 3042 AP cannot issue I/O instructions, it can process data read and to be written by the 3033 Processor. Therefore, in effect, the I/O devices in a 3033 attached processor configuration are pooled for use by both processors. More than half the total number of I/O devices present can be allocated to an individual job step when necessary. The pooling of I/O devices and sharing of processor storage permits the execution of jobs with larger processor storage and I/O device requirements than can be handled using one system in a configuration with two uncoupled systems.

The sharing of processor storage and the ability of the 3033 Processor to access all I/O devices in the configuration also enables the OS/VS2 MVS control program to automatically handle peak load

situations within jobs and to balance the processing across the two processors. Manual balancing of the workload between two systems, as is required for two uncoupled systems, is not required for a tightly coupled configuration.

Through pooling, the number of I/O devices in a tightly coupled multiprocessing configuration can be less than the number of I/O devices needed for two uncoupled systems that are to handle the same large I/O job or peak load direct access storage requirements.

Since there is only one control program for an attached processor configuration, there is no need to split any data base into two parts, one for each system, or to use Shared DASD support in order to share a data base between the two systems, as is required for two uncoupled systems. The use of Shared DASD support results in reduced throughput for two uncoupled systems because of the interference it introduces. This throughput reduction is not incurred in an attached processor configuration since there is only one OS/VS2 MVS control program and it can maintain the integrity of a shared data base without using Shared DASD support.

SECTION 50: PROGRAMMING SYSTEMS SUPPORT

50:05 GENERAL DESCRIPTION

Basic support of the 3033 Processor Complex is provided by OS/VS2 MVS as of Release 3.7, OS/VS2 SVS Release 1.7, OS/VS1 as of Release 6, VM/370 as of Release 5, and ACP (Airlines Control Program) Version 9 Release 2. The MVT/3031/3032/3033 programming RPQ is provided to enable OS MVT Release 21.8 to operate on the 3033 Processor. This support is designed primarily to aid in the transition from MVT to OS/VS2 MVS.

The basic support provided by these programming systems includes (1) recognition of the 3033 Processor identification stored by the STORE CPU ID instruction and support of processor-dependent control program initialization as required, (2) support of the 3036 Console, (3) recognition of the 3033 channels and I/O extended logout area of 576 bytes, and (4) EREP support of the service record file and the frame verbiage approach to formatting processor-dependent logout data of the 3033 Processor.

OS/VS2 MVS also recognizes the two additional channel error conditions (interface inoperative and channel-not-operational) and provides error recovery procedures that include use of the CLRCH instruction. OS/VS2 SVS, OS/VS1, and VM/370 do not attempt recovery using the CLRCH instruction after these two channel errors.

The common segment facility, low-address protection, IPTE and TSPT instructions, and MVS-dependent instructions of the System/370 Extended Facility are supported only by the MVS/System Extensions program product. The virtual machine extended facility assist of the System/370 Extended Facility is supported by the VM/System Extensions program product.

The 3033 Multiprocessor Complex operating in multiprocessor mode is supported by OS/VS2 MVS with appropriate selectable units installed. The channel set switching function is supported for recovery processing that occurs after an uncorrectable processor failure.

The 3033 Attached Processor Complex operating in attached processor mode is supported by OS/VS2 MVS and VM/370. Channel set switching is supported by both of these programming systems.

The MVS/System Extensions, VM/System Extensions, and VM/Basic System Extensions program products also support the 3033 Processor Complex, Attached Processor Complex, and Multiprocessor Complex. These program products offer additional performance improvements.

50:10 MVS/SYSTEM EXTENSIONS PROGRAM PRODUCT

The MVS/System Extensions program product is designed to improve the performance of MVS. It requires the System/370 Extended Facility (standard on 3033, 3032, and 3031 Processors) or the System/370 Extended feature (optional on Models 158 and 168) and OS/VS2 MVS. This program product supports multiprocessor mode operations for the 3033 Multiprocessor Complex, attached processor mode operations for a 3033 or 3031 Attached Processor Complex, and attached processor and multiprocessor mode operations for Models 158 and 168, as well as uniprocessor mode of operation.

Via its support of the System/370 Extended Facility (or System/370 Extended feature), the MVS/System Extensions program product provides the following:

- A reduction in the processor time required to execute certain frequently used control program functions via the use of hardware implemented assists.
- A reduction in the time required for address translation through more efficient use of the translation lookaside buffer and a reduction in buffer storage interference caused by the address translation process.
- Improved system availability via support of a new level of storage protection for lowest addressed processor storage that contains data vital to continued system operation.

The MVS/System Extensions program product provides the following improvements via programming:

- A reduction in the processor time required to execute certain frequently used control program functions through a reduction in path length for those functions and a reduction in programmed contention.
- Improved installation control over resource utilization via changes to the System Resource Manager and Input/Output Supervisor and as a result of improved resource usage recording.
- Improved performance for attached processor and multiprocessor mode operations via a reduction in interprocessor interference.

User-written programs that operate under MVS and use standard external interfaces will operate under MVS/System Extensions without modification. Programs that interrogate control program control blocks may require modification. MF/1 and DSS are not supported by this program product. For additional information regarding MVS/System Extensions, see OS/VS2 MVS System Extensions General Information Manual (GC28-0872).

50:15 VM/SYSTEM EXTENSIONS AND VM/BASIC SYSTEM EXTENSIONS PROGRAM PRODUCTS

The VM/System Extensions and VM/Basic System Extensions program products can be utilized in 3033, 3032, and 3031 Processors and Models 135, 138, 145, 148, 155 II, 158, 165 II, and 168, as well as 4300 Processors. Attached processor mode operations in a 3033 or 3031 Attached Processor Complex and Model 158 or 168 attached processor configuration are also supported.

The VM/System Extensions program product is designed to improve VM/370 performance and provide additional function. It provides the following:

- Improved performance for all operating systems that execute in a virtual machine by providing the functions of the no longer available VM/370 Resource Management PRPQ, which offers throughput improvements, improved response to trivial commands, and additional installation management support.
- Additional performance gains for DOS/VS, OS/VS1, and OS/VS2 (SVS and MVS) operating systems executing in a virtual machine by more efficient management of the shadow page and shadow segment tables that are required for virtual storage operating systems.

- Support of the MVS/System Extensions program product. An MVS operating system with the MVS/System Extensions program product installed can execute in a virtual machine when VM/System Extensions is utilized. This enables the System/370 Extended Facility for 3031, 3032, and 3033 Processors and the System/370 Extended feature for Models 158 and 168, plus the other facilities of MVS/System Extensions to be utilized in a VM/370 environment during uniprocessor or attached processor mode operations.
- Virtual Storage preservation support. This facility enables the contents of virtual storage in user-specified virtual machines (defined during system generation) to be saved on direct access storage whenever the virtual machine is abnormally terminated. It also provides a means of initiating recovery procedures for the operating system executing in the virtual machine at the time of the abnormal termination.
- Spooling of accounting records to direct access storage instead of punching accounting cards.
- Writing spool files contained on disk to tape, instead of transcribing them to a printer or punch, for example.
- CMS processing of standard tape labels. This support includes CMS checking of standard labels on input tapes, writing standard labels on output tapes, a user exit for processing nonstandard labels on tapes during execution of CMS macro simulations and some CMS tape commands, and support of user-written routines that process standard user labels during DOS/VS and OS macro simulation under CMS.
- Additional support for 3277 and 3278 Display Stations. Using the DIAGNOSE interface, a virtual machine can directly manage the screen of a 3277 or 3278.
- Additional support for 3278 display stations. Alternate screen sizes (2560 and 3440 characters in addition to 1920), the optional additional 12 program function keys (for a total of 24), APL keyboards, and six new national usage characters are supported.

The VM/System Extensions program product is recommended for MVS, SVS, and VS1 installations that use VM/370 for testing, conversion, timesharing or production operations. It is also recommended for installations that use the 2305 Model 2 Fixed Head Storage Device for paging or that heavily use CMS.

The VM/Basic System Extensions program product provides support of a subset of the facilities offered by the VM/System Extensions program product. VM/Basic System Extensions is recommended for VS1, DOS/VSE, and DOS/VS installations that utilize VM/370 for interactive processing and testing (as provided by CMS), conversion, and production operations.

SECTION 60: HIGHLIGHTS OF THE 3033 PROCESSOR COMPLEX FOR MODEL 168 USERS

The 3033 Processor Complex offers Model 168 users significantly increased price performance (as a result of several improvements detailed in this section); reductions in space, cooling, and power requirements (through use of new technology and as a result of the channel group design); and additional availability and serviceability features provided by the new design of the 3036 Console.

PRICE PERFORMANCE IMPROVEMENTS

More Standard Features

The System/370 Extended Facility, high-speed multiply capability, twelve channels, and channel indirect data addressing are standard in the 3033 Processor and optional for the 3168 Processing Unit (Model 1 or 3). The minimum processor storage size is four megabytes, instead of one, and standard buffer size is 64K instead of 8K or 32K.

Faster Processor Cycle

The 57-nanosecond processor cycle of the 3033 Processor, versus the 80-nanosecond cycle of the 3168 Processing Unit, results in faster buffer fetches, reduced processor storage cycle times, and faster instruction execution times.

Improvements in IPPF and Execution Function Design

The major changes in the IPPF and execution function (detailed in Section 10:10) result in faster instruction execution.

Processor Storage Control Function Changes

More requests for processor storage are generated by the IPPF than by any other component (execution function, channels, or console). Conflicts between the IPPF and PSCF for access to processor storage that occur in the Model 168 are reduced in the 3033 Processor to improve internal performance by the following:

- An outstanding fetch request is given priority over the outstanding store requests (the opposite is true in the 3168 processor).
- The highest priority requester in the IPPF for a given logical storage element is determined by the IPPF, instead of by the PSCF as in the 3168 Processing Unit.

High Speed Buffer Improvements

- More buffer storage is provided (64K versus 32K for the 3168-3 and 16K for the 3168-1) to improve the buffer hit ratio.
- A buffer storage block is 64 bytes (eight doublewords), instead of 32 bytes (four doublewords) as in the Model 168 buffer. This block size matches the eight-way interleaving of processor storage in the

3033 Processor and makes more data available in the buffer after a block load, which improves sequential instruction execution and sequential data access times.

- A doubleword fetch from the buffer requires 114 nanoseconds, versus 160 nanoseconds in the Model 168.
- A row deletion causes a smaller percentage of the total buffer to be deleted (6 percent versus 12).

Faster Dynamic Address Translation

Several improvements in the TLB and STO-stack have been made:

- The TLB is oriented to a 4K page size, instead of a 2K page size as in the TLB in the Model 168. Thus, when 4K pages are used, only one TLB entry is required, versus two TLB entries in the Model 168 TLB. Therefore, 128 translations, instead of 64, can be maintained in the TLB for 4K pages (representing 512K of processor storage instead of 256K).
- The replacement algorithm (which determines which entry of a pair at a given one of the 64 TLB locations is assigned) is least recently used, instead of first-in, first-out as in the Model 168 TLB.
- The time required to purge the entire TLB is 16 processor cycles, instead of 64 as for the Model 168 TLB. Less full purging is required when the IPTE is utilized (optional facility in the Model 168).
- The STO-stack contains 29 entries, instead of 6 as in the STO-stack for the Model 168, and a different replacement algorithm is implemented. The use of 29 entries enables translations from 23 more virtual storages to be kept in the TLB at the same time than in the Model 168 without the necessity of partial TLB purge operations, which are required whenever a new virtual storage is added to the 6-entry STO-stack in the Model 168. In the 3033 Processor, there is no partial purging for the addition of a new virtual storage. When the 29-entry STO-stack is full, the entire TLB is purged.

The changes made in the TLB and STO-stack in the 3033 Processor reduce the amount of TLB purging required and reduce the time to perform a full TLB purge when it is necessary.

Processor Storage Improvements

Twice the amount of processor storage is available for a 3033 Processor (16M-bytes versus 8M-bytes) and the access to processor storage has been improved:

- Processor storage is eight-way doubleword interleaved, instead of four-way doubleword as in the Model 168, to permit twice the amount of parallel access.
- Read and write cycle times are reduced from 320 and 640 nanoseconds to 285 and 456 nanoseconds and access times are reduced from 400 and 640 nanoseconds to 285 and 456 nanoseconds. The aggregate data rate of which 3033 processor storage is capable is 138 MB/sec compared with 100 MB/sec for 3168 processor storage.

CHANNEL ENHANCEMENTS

The implementation of channel groups that are contained within the 3033 Processor but controlled by independent microprocessors offers several advantages without losing the advantage of less instruction processor interference offered by standalone versus integrated channels:

- Less space, power (approximately 30 percent), and cooling (approximately 20 percent less for combined water and air) are required for a 3033 Processor Complex than for a Model 168 with the same number of channels and twice the maximum amount of processor storage. The space reduction becomes more significant as the number of channels in the Model 168 increases.
- Since a maximum of 16 channels, instead of 12, are available for the 3033 Processor and twelve are standard, the I/O device configuration can be spread across the available channels to minimize channel overloads and more paths to critical devices can be provided via switching features. Three byte multiplexer channels, instead of a maximum of two, as for the Model 168, are available.
- A malfunctioning channel group can be serviced (locally or remotely) concurrently with normal system operation, assuming enough operable channels are available in the other channel group(s) for critical devices. This is not possible in a Model 168 configuration which requires dedication of the processor when diagnostics are executed on a standalone channel.
- Channel groups can be powered up and down separately from each other and other processor components. This enables certain fixes to be applied without interference with normal system operation.
- More subchannels are provided for the channels of a 3033 Processor (256 nonshared or up to eight shared with reduced nonshared per each byte and block multiplexer channel). The 2870 Multiplexer Channel can have a maximum of 192 subchannels (plus four selector subchannels) while the 2880 Block Multiplexer Channel always has only one shared subchannel and 56 (or 256 with an option) nonshared subchannels.
- The implementation of significantly more channel logout data, the two additional channel error conditions, the capability of faster engineering change installation (because of microcode instead of hardware control), and the CLRCH instruction for channel recovery improve the serviceability and availability of the channels, as compared with Model 168 channels.

AVAILABILITY AND SERVICEABILITY FEATURES

The 3036 Console and the channel group implementation are designed to improve the availability and serviceability characteristics of a 3033 Processor as compared with those of a Model 168 Model 3, which includes a service processor that is not available in the Model 168 Model 1, as follows:

- Channel availability and serviceability is improved as discussed previously under "Channel Enhancements".
- The dual station design of the 3036 provides console device backup and interchangeability of console functions (operator and service). The 3066 has only one CRT display and its functions are separate from those of the indicator viewer.

- The customer engineer can use a CRT display and keyboard to perform service operations concurrently with normal processing, as a result of the dual station design of the 3036 Console. The CE panel must be used to control concurrent maintenance operations involving the service processor in a Model 168 Model 3 configuration.
- The remote service options for the 3036 Console include allowing a remote specialist to control operation of the concurrent maintenance functions. This option is not provided for a Model 168 Model 3.
- More concurrent maintenance operations can be performed using the 3036 Console than the service processor of a Model 168 Model 3 and when the processor is dedicated to maintenance functions, two components can be serviced concurrently.
- Voltage monitoring enables transient power faults to be detected and located more easily.
- Diagnostic testing can be done faster because of the building block approach for fault locating tests that is implemented in the 3033 Processor as a result of the new 3036 Console design.

NEW TECHNOLOGY

As indicated in Section 05, new logic technology that provides more circuits per logic chip is implemented in the 3033 Processor. This results in speed improvements (such as the 57-nanosecond versus 80-nanosecond cycle time) as well as physical space savings. In addition, the 4K-bit processor storage chip used in certain processor storage models of the 3033 Processor provides twice the amount of processor storage in the same space required by the 2K-bit chip.

SECTION 70: COMPARISON TABLE

This table has been included for quick reference. It compares hardware features of System/370 Models 165, 165 II, and 168 (Models 1 and 3), and the 3033 Processor Complex.

70:05 COMPARISON TABLE OF HARDWARE FEATURES OF SYSTEM/370 MODELS 165, 165 II, AND 168 (MODELS 1 AND 3), AND THE 3033 PROCESSOR COMPLEX

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
I. Processor				
A. BC mode of system operation	Standard	Standard	Standard	Standard
B. EC mode of system operation	Not implemented	Standard	Standard	Standard
C. Instruction set				
1. Standard set (binary arithmetic)	Standard	Standard	Standard	Standard
2. Decimal arithmetic	Standard	Standard	Standard	Standard
3. Floating-point arithmetic	Standard	Standard	Standard	Standard
4. Extended precision floating-point	Standard	Standard	Standard	Standard
5. New instructions listed in Section 10:05	Standard except for CLEAR I/O, MONITOR CALL, conditional swapping, PSW key handling, clock comparator and CPU timer, system mask, IPTE, CLRCH, TSPT, MVS-dependent, and dynamic address translation instructions, which are not available.	Standard except IPTE, TSPT, and CLRCH, which are not available.	Standard except IPTE, TSPT, CLRCH, and MVS-dependent instructions. IPTE, TSPT, and MVS-dependent instructions are included in the optional. System/370 Extended feature. CLRCH is not provided.	Standard
D. Overlap of instruction fetching and preparation with instruction execution	Instruction unit can prepare up to three instructions while execution unit executes one instruction. Imprecise interruptions can occur.	Same as Model 165 except imprecise interruptions cannot occur.	Instruction unit can prepare up to four instructions while one is executed. Instruction and execution unit implementation is enhanced over Model 165 and imprecise interruptions cannot occur.	Same as Model 168 except that implementation is enhanced over Model 168 (many more instructions can be prefetched, etc.)

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
E. High-speed multiply	Optional	Optional	Optional	Basic to implementation of execution function
F. Processor cycle time	80 nanoseconds 8-byte data path	Same as Model 165	Same as Model 165	57 nanoseconds 8-byte data path
G. Dynamic address translation	Not available	Standard (with 128-entry TLB and 6-entry STO-stack)	Same as Model 165 II	Standard (with 128-entry TLB and 29-entry STO-stack) plus performance enhancements, such as orientation to 4K instead of 2K page size
H. Interval timer	Standard (3.33 ms resolution)	Standard (3.33 ms resolution)	Standard (3.33 ms resolution)	Standard (3.33 ms resolution)
I. Time-of-day clock	Standard	Standard	Standard	Standard
J. CPU timer and clock comparator	Not available	Standard	Standard	Standard
K. Monitoring feature	Not available	Standard	Standard	Standard
L. Program event recording	Not available	Standard	Standard	Standard
M. Direct control	Standard	Standard	Standard	Standard
N. Interruption for SSM instruction	Not implemented	Standard	Standard	Standard
O. System/370 Extended Facility	Not available	Not available	Optional System/370 extended feature includes same items	Standard
P. Compatibility features (all are optional and mutually exclusive except where noted otherwise)	1. 7070/7074 2. 7080 (for both 705 and 7080) 3. 709/7090/7094 7094II (does not include 704, 7040, 7044)	Same as Model 165	Same as Model 165	None available
Q. Control logic	Capacitor ROS and monolithic WCS	Same as Model 168	Monolithic ROS and monolithic WCS	Monolithic RCS
R. Instruction retry by hardware	Yes	Yes	Yes	Yes

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
S. Machine check interruption	Occurs after corrected and uncorrected errors. There are four types of machine check interruptions and many are individually maskable	Occurs after corrected and uncorrected errors. There are seven types and many are individually maskable	Same as Model 165 II	Same as Model 165 II plus external damage indicated for a hard channel error.
T. Fixed storage area size in lower storage (including logout area for machine and channel errors)	1504 bytes reducible to 512 if extended logout area of 992 bytes is moved	Same as Model 168	1928 bytes reducible to 512 if the extended logout area of 1416 bytes is moved	Same as Model 168
U. Multiprocessor systems	<ol style="list-style-type: none"> 1. A multisystem feature is not available 2. A Model 165 can be a support or a main processor in an ASP configuration 	Same as Model 165	<ol style="list-style-type: none"> 1. The 3068 Multisystem Communication Unit is used to connect two Model 168 systems (any combination of Models 1 and 3) together for a tightly coupled multiprocessing configuration 2. JES2 and JES3 (Job Entry Subsystem) of OS/VS2 MVS support the Model 168 in loosely coupled multiprocessing configurations 3. A Model 168 can be a main or support processor in an ASP configuration 	<ol style="list-style-type: none"> 1. The 3033 Multiprocessor Complex is a tightly coupled multiprocessing configuration 2. Same as Model 168 3. Same as Model 168

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
V. Attached Processor System	Not available	Not available	Supported (RPQs are required for a Model 1)	Available
W. Power Warning	Not available	Not available	Optional	RPQ feature (operation can be enabled or disabled using the configuration frame)
X. Virtual Machine Assist	Not available	RPQ feature	RPQ feature	RPQ feature
Y. Remote analysis unit	Yes - 2955 Remote Analysis Unit is optional.	Same as Model 165	Yes - 2955 Remote Analysis Unit is optional in the Model 1, service processor is standard in the Model 3.	Yes (2955 operation is simulated and one console processor can be used for remote service operations).
Z. Integrated Storage Controls	Not available	Not available	Optional for attachment of 3330-series, 3340/3444, and/or 3350 disk storage, or the 3850 Mass Storage System without a 3830	Not available
II. Storage				
A. Processor (main) storage sizes	512K 1024K 1536K 2048K 3072K	1024K 2048K 3072K	1024K 2048K 3072K 4096K 5120K 6144K 7168K 8192K	4096K 8192K 12,288K 16,384K
B. Type of processor storage	Ferrite cores	Ferrite cores	Monolithic technology	Monolithic technology
C. Processor storage interleaving	Storage is 4-way doubleword interleaved	Same as Model 165	Same as Model 165	Storage is 8-way doubleword interleaved
D. High-speed buffer storage	8K is standard, 8K more can be added. 80-nanosecond cycle.	Same as Model 165	Model 1 is same as Model 165. 32K is standard in the Model 3. Buffer row and block deletion are implemented in both models.	64K is standard. 57-nanosecond cycle. Buffer row and block deletion are standard.

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
E. Processor storage validity checking	ECC checking on a doubleword. All single-bit errors are corrected and all double-bit and some multiple-bit errors are detected by hardware.	Same as Model 165	Same as Model 165	Same as Model 165 except improved code detects more multiple-bit errors
F. Byte-oriented operands	Standard	Standard	Standard	Standard (significantly less degradation occurs than for Models 165, 165 II, and 168)
G. Store and fetch protection	Standard	Standard	Standard	Standard
H. Processor storage cycle times	Read/write cycle is 2 microseconds for 32 bytes	Same as Model 165	Read/write cycle for a doubleword is 320 nanoseconds	Read/write cycle is 285 nanoseconds for a doubleword
III. Channels				
A. Total number per system	1. Up to 7 standard 2. Up to 12 with Extended Channels optional feature	Same as Model 165	Same as Model 165	1. Twelve standard 2. Four additional are optional
B. 2870 Multiplexer Channel	One or two can be attached (192 subchannels). Maximum byte mode rate of 110 KB/sec, which is reduced when selector subchannels are installed.	Same as Model 165	Same as Model 165	Does not attach. Channels 1 and 6 are standard (and channel 12 is an optional) byte multiplexer channels. (256 subchannels of which up to eight can be designated as shared. Maximum byte mode rate of up to 40 KB/sec to 75 KB/sec)
C. 2860 Selector Channel (1.3 MB/sec.)	A maximum of six can be attached	Same as Model 165	Same as Model 165	Does not attach. Selector mode standard for all block multiplexer channels

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
D. 2880 Block Multiplexer Channel (1.5 MB/sec). Two-Byte Interface feature permits a 3.0 MB/sec. data rate	A maximum of 6 can be attached without the Extended Channels feature, a maximum of 11 with this feature.	Same as Model 165	Same as Model 165	Does not attach. Standard block multiplexer channels are 1 to 5 and 7 to 11. Channels 12 to 15 or 13 to 15 are optional. Maximum data rate is up to 1.5 MB/sec. Two-Byte Interface (optional on channels 1, 7, and 12 or 13) provides up to 3.0 MB/sec maximum.
1. Maximum number of subchannels	1 shared and 56 or 64 nonshared without extended unit control words feature, 1 shared and up to 256 nonshared with feature	Same as Model 165	Same as Model 165	256 per channel of which up to 8 can be designated as shared with the rest nonshared
E. Extended Unit Control Words on the 2880	-	Optional	Optional	Not available
F. Channel retry data provided after channel error	Yes	Yes (I/O extended logout only)	Same as Model 165 II	Yes (limited channel logout)
G. Additional channel error logging and recovery	No	No	No	Yes (extended channel logout and two additional errors reported and CLRCH instruction provided for recovery)
H. Channel-to-Channel Adapter	Optional on 2860 (RPQ required for installation on 2880)	Optional on 2860 (RPQ required for installation on 2880)	Optional on 2860 (RPQ required for installation on 2880)	Optional on one byte or block multiplexer in each standard channel group (two maximum per processor)
I. Channel indirect data addressing	Not available	Optional (required by the virtual storage programming systems)	Optional (required by the virtual storage programming systems)	Standard

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
IV. Operator console devices	<p>1. Standalone 3066 Model 1 System Console is required. It includes:</p> <ul style="list-style-type: none"> a. A CRT-key-board combination for operator/system communication b. An indicator viewer c. A microfiche document viewer d. A processor storage configuration plug-board e. A system activity meter f. A device for loading WCS and microdiagnostics <p>The store status function is not provided.</p> <p>2. 2150 Console with Operator Console Panel and 1052-7 Printer-Keyboard is optional.</p> <p>3. Optionally, other devices can be used as secondary consoles.</p>	<p>Same as Model 165 except store status is provided.</p>	<p>1. Standalone 3066 Model 2 System Console provides same features as 3066 Model 1 and store status function. Other consoles can be attached as for Model 165.</p>	<p>1. Standalone 3036 Console with dual station design provides two CRT displays and keyboards, two console processors, two diskette drives, two standard I/O interfaces, and one modem. Store status recording and system status recording are provided. Concurrent maintenance functions not available for the Model 168 are provided.</p> <p>2. 2150 Console can be attached but Remote Operator Console Panel feature is not supported.</p> <p>3. Optionally, other devices can be used as secondary consoles.</p>
V. I/O devices				
A. 3505 Card Reader 3525 Card Punch	Yes	Yes	Yes	Yes
B. 3211 Printer	Yes	Yes	Yes	Yes
C. 3803/3420 Magnetic Tape Subsystem (Models 3,5,7 and 4,6,8)	Yes	Yes	Yes	Yes

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
D. Direct access devices (2311, 2314/2319, 2303, 2301, and 2321)	Yes	Yes	Yes	All attach except 2301, 2303, 2321
E. 3330-series disk storage	Yes (RPQ required for Model 11)	Yes (all models)	Yes (all models)	Yes (all models)
1. 3830 Storage Control Model 1	Yes	Yes	Yes	Yes
2. 3830 Storage Control Model 2	Yes	Yes	Yes	Yes
3. Integrated Storage Controls feature	No	No	Yes	No
F. 2305 facility Models 1 and 2	Yes on 2880	Yes on 2880	Yes on 2880	Yes
G. 3340 direct access storage facility	No	Yes (attachment via 3830 Model 2)	Yes (attachment via 3830 Model 2 and Integrated Storage Controls)	Same as Model 165 II
H. 3344 Direct Access Storage	No	Yes (attachment via 3830 Model 2)	Yes (attachment via 3830 Model 2 and Integrated Storage Controls)	Same as Model 165 II
I. 3350 Direct Access Storage	RPQ required (attachment via 3830 Model 2)	Yes (attachment via 3830 Model 2)	Yes (attachment via 3830 Model 2 and Integrated Storage Controls)	Same as Model 165 II
J. 3410/3411 Magnetic Tape Subsystem	No	No	No	No
K. 3540 Diskette I/O Unit	No	Yes	Yes	Yes
L. 3600 Finance Communication System	No	Yes	Yes	Yes
M. 3650 Retail Store System	No	Yes	Yes	Yes
N. 3660 Supermarket System	No	Yes	Yes	Yes
O. 3704, 3705-I, and 3705-II Communications Controllers	Yes	Yes	Yes	Yes
P. 3740 Data Entry System	Yes	Yes	Yes	Yes

<u>Hardware Feature</u>	<u>Model 165</u>	<u>Model 165 II</u>	<u>Model 168 (Models 1 and 3)</u>	<u>3033 Processor Complex</u>
Q. 3767 Communication Terminal	Yes	Yes	Yes	Yes
R. 3770 Data Communication System	Yes	Yes	Yes	Yes
S. 3780 Data Communications Terminal	Yes	Yes	Yes	Yes
T. 3790 Communication System	No	Yes	Yes	Yes
U. 3800 Printing Subsystem	No	Yes	Yes	Yes
V. 3838 Array Processor	No	No	Yes	Yes
W. 3850 Mass Storage System	No	Yes via 3830 Model 3	Yes via 3830 Model 3 and Integrated Storage Controls	Yes via 3830 Model 3
X. 3881 Optical Mark Reader	No	No	No	No
Y. 3886 Optical Character Reader	No	Yes	Yes	Yes
Z. 3890 Document Processor	No	Yes	Yes	Yes
AA. 3895 Deposit Processing System	No	Yes	Yes	Yes
BB. 3270 Information Display System	Yes	Yes	Yes	Yes
CC. 3250 Graphics Display System	No	Yes	Yes	Yes

70:10 OS/VS SUPPORT OF THE 3033 PROCESSOR COMPLEX

<u>Hardware Feature</u>	<u>OS/VS1</u>	<u>OS/VS2 SVS</u>	<u>OS/VS2 MVS</u>
I. PROCESSOR			
A. Mode of system operation	EC and DAT modes only. One virtual storage of up to 16 million bytes is supported. Up to 52 partitions of which 15 can be problem program.	EC and DAT modes only. One virtual storage of 16 million bytes is supported. Up to 63 problem program regions of which up to 42 can be TSO foreground regions.	EC and DAT modes only. Multiple virtual storages are supported. Each user has one 16-million-byte virtual storage for user programs, system programs, shared data, and shared program areas. The maximum number of concurrent users is limited only by the availability of system resources (external page and real storage).
B. Instruction set			
1. Standard set (binary arithmetic)	All languages	All languages	All languages
2. Decimal arithmetic	All languages except FORTRAN	All languages except FORTRAN	All languages except FORTRAN
3. Floating-point arithmetic	All languages except RPG	All languages except RPG	All languages except RPG
4. Extended precision floating-point	Assemblers F and H, PL/I Optimizing Compiler, PL/I Checkout Compiler, FORTRAN H, FORTRAN H-Extended	Same as OS/VS1	Same as OS/VS1
5. New instructions	All are supported by the System Assembler	Same as OS/VS1	Same as OS/VS1
C. Interval timer	Supported for all timing facilities (except time of day) unless the extended timer option is included in the VS1 control program	Not supported	Not supported
D. Time-of-day clock	Supported for time of day	Same as OS/VS1	Same as OS/VS1

<u>Hardware Feature</u>	<u>OS/VS1</u>	<u>OS/VS2 SVS</u>	<u>OS/VS2 MVS</u>
E. Clock comparator and CPU timer	Supported for job step and interval timing when extended timer option is included in the VS1 control program	Supported for timing facilities except for time of day	Supported for timing facilities except time of day
F. Expanded machine check interruptions	Supported by MCH	Same as OS/VS1	Same as OS/VS1
G. Monitoring feature	Supported by GTF and an Assembler mnemonic	Same as OS/VS1	Same as OS/VS1
H. Program event recording	Not supported for 3033 Processor	Not supported for 3033 Processor	Not supported for 3033 Processor
I. Interruption for SSM instruction	Supported	Supported	Supported
J. System/370 Extended Facility	Not supported	Not supported	Supported by MVS/System Extensions program product using an OS/VS2 MVS Release 3.7 base
K. Multiprocessor systems	Not supported	Tightly coupled multiprocessing is not supported. Loosely coupled multiprocessing is supported via ASP.	The 3033 Multiprocessor Complex (including Channel Set Switching) is supported. Loosely coupled multiprocessing is supported by JES2, JES3, and ASP.
II. STORAGE			
A. Real storage sizes	All are supported	All are supported	All are supported
B. Byte-oriented operands	Programmers can use the byte alignment hardware facility in Assembler programs	Same as OS/VS1	Same as OS/VS1
C. Store and fetch protection	Store and fetch protection are supported for all partitions	Store and fetch protection are supported for all regions	Store and fetch protection are supported for all virtual storages
III. CHANNELS			
A. Byte multiplexer channels	One or two are supported	One or two are supported	Up to three are supported
B. Block multiplexer and selector channels	Supported	Supported	Supported
C. Channel retry performed	Yes	Yes	Yes
D. Channel indirect data addressing	Supported	Supported	Supported
E. Channel recovery using the CLRCH instruction	Not supported	Not supported	Supported

<u>Hardware Feature</u>	<u>OS/VS1</u>	<u>OS/VS2 SVS</u>	<u>OS/VS2 MVS</u>
IV. CONSOLES			
A. 3036 Console	Supported	Supported	Supported
B. Alternate and additional consoles supported	Yes	Yes	Yes
V. I/O DEVICES			
A. 3505 Card Reader and 3525 Card Punch	Supported	Supported	Supported
B. 3211 Printer	Supported	Supported	Supported
C. 3803/3420 Magnetic Tape Subsystem (Models 3, 5, 7 and 4, 6, 8)	Supported	Supported	Supported
D. 2314/2319 facilities	Supported for system residence, data sets, paging devices, JES spooling devices and data sets, and SYSIN devices. Record Overflow and channel switching features are supported.	Supported for system residence, data sets, paging devices, SYSIN and SYSOUT data sets, and SYSIN devices. Record Overflow and channel switching features are supported.	Same as OS/VS1
E. 3330-series Direct Access Storage	Same as V.D. above. RPS, multiple requesting, sixteen-drive addressing, 32 Drive Expansion, Two-Channel Switch, Additional, 3333 String Switch, and Record Overflow are supported. All models are supported.	Same as V.D. above. RPS, multiple requesting, sixteen-drive addressing, 32 Drive Expansion, Two-Channel Switch, Additional, 3333 String Switch, and Record Overflow are supported. All models are supported.	Same as OS/VS1
F. 3340 and 3344 Direct Access Storage	Support same as for 3330-series	Support same as for 3330-series	Support same as for 3330-series

<u>Hardware Feature</u>	<u>OS/VS1</u>	<u>OS/VS2 SVS</u>	<u>OS/VS2 MVS</u>
G. 3350 Direct Access Storage	Supported in native and 3330 compatibility modes as for 3330-series	Supported in native and 3330 compatibility modes as for 3330-series	Supported in native and 3330 compatibility modes as for 3330-series
H. 2305 Facility Models 1 and 2	Same as V.D. above except for SYSIN devices. RPS and multiple requesting are supported.	Same as V.D. above. RPS and multiple requesting are supported.	Same as OS/VS1
I. 3540 Diskette I/O Unit	Supported as a SYSIN and SYSOUT device (not an I/O device)	Same as OS/VS1	Same as OS/VS1
J. 3600 Finance Communication System	Supported attached to a 3704/3705 in NCP/VS mode by VTAM and TCAM through VTAM	Supported attached to a 3704/3705 in NCP/VS mode by VTAM	Same as OS/VS1
K. 3650 Retail Store System	Supported in binary synchronous control mode attached to a 3704/3705 in emulation mode by BTAM. Supported in synchronous data link control mode attached to a 3704/3705 in NCP/VS mode by VTAM and TCAM through VTAM	Not supported	Supported in synchronous data link control mode attached to a 3704/3705 in NCP/VS mode by VTAM and TCAM through VTAM
L. 3660 Supermarket System	Supported in binary synchronous mode attached to a 3704/3705 in emulation mode by BTAM. Supported in synchronous data link control mode attached to a 3704/3705 in NCP/VS mode by VTAM.	Supported in binary synchronous mode attached to a 3704/3705 in emulation mode by BTAM	Supported in synchronous data link control mode attached to a 3704/3705 in NCP/VS mode by VTAM
M. 3704 and 3705 Communications Controllers	Supported in emulation mode Supported in NCP mode by TCAM Supported in NCP/VS mode by TCAM and VTAM	Supported in emulation mode Supported in NCP mode by TCAM Supported in NCP/VS mode by TCAM and VTAM	Same as OS/VS1

<u>Hardware Feature</u>	<u>OS/VS1</u>	<u>OS/VS2 SVS</u>	<u>OS/VS2 MVS</u>
N. 3740 Data Entry System	Supported (BTAM, TCAM, and TCAM through VTAM)	Supported (BTAM,TCAM, and as a System/3 by VTAM)	Same as OS/VS1
O. 3767 Communication Terminal	Supported (as a start/stop device) attached to a 3704/3705 in emulation mode by BTAM and TCAM. Supported attached to a 3704/3705 in NCP/VS mode by VTAM and TCAM through VTAM.	Supported (as a start stop device) attached to a 3704/3705 in emulation mode by BTAM and TCAM. Supported attached to a 3704/3705 in NCP/VS mode by VTAM.	Same as OS/VS1
P. 3770 Data Communication System	Supported for synchronous data link control (SDLC) operations attached to a 3704/3705 in NCP/VS mode by VTAM and TCAM through VTAM. Supported for binary synchronous communications (BSC) operations to a 2701 or 3704/3705 by 2770 support in BTAM, TCAM, and VTAM.	Supported for binary synchronous communication operations attached to a 2701 or 3704/3705 by 2770 support in BTAM and TCAM. Supported for synchronous data link control operations attached to a 3704/3705 in NCP/VS mode by VTAM.	Same as OS/VS1
Q. 3780 Data Communications Terminal	Supported (BTAM, TCAM, and TCAM via VTAM)	Supported (BTAM, TCAM and VTAM)	Same as OS/VS1
R. 3790 Communication System	Supported attached to a 3704/3705 in NCP/VS mode by VTAM	Same as OS/VS1	Same as OS/VS1
S. 3800 Printing Subsystem	Supported	Supported	Supported
T. 3838 Array Processor	Supported	Not supported	Supported
U. 3850 Mass Storage System	Supported	Supported	Supported
V. 3886 Optical Character Reader	Supported	Not supported	Supported
W. 3890 Document Processor	Supported	Supported	Supported
X. 3895 Deposit Processing System	Supported	Supported	Supported
Y. 3270 Information Display System	Supported	Supported	Supported
Z. 3250 Graphics Display System	Supported	Supported	Supported

INDEX (Sections 01 through 70)

- architecture design 9
- attached processor complex 108
 - advantages 121
 - channel set switching 111, 117
 - channels 111
 - configuration components 108
 - configuration frame 116
 - differences between uniprocessor and attached processor hardware 119
 - instructions 119
 - interprocessor hardware communication 118, 119
 - interprocessor programmed communication 119
 - I/O devices 111
 - modes of operation 109, 116
 - performance 108
 - power control 111
 - prefixing 119
 - processor
 - addressing 119
 - features 110
 - storage 110
 - programming support 60
 - space requirement 108
 - time-of-day clock 119
 - 3033 Model A Processor 110
 - 3036 Console 114
 - 3038 Multiprocessor Communication Unit 113
 - 3042 Attached Processor 113
- basic control mode 9
- block, buffer 29
- block multiplexer channels
 - addresses 37
 - Channel-to-Channel Adapter 37
 - data rate 37
 - maximum number 37
 - priorities 41
 - selector mode 37
 - subchannels 39
 - Two-Byte Interface feature 37
- buffer storage
 - access time 27
 - activity lists 29
 - address array 29
 - block definition 29
 - block deletion 31
 - capacity 31
 - column definition 29
 - cycle time 27
 - disabling 28
 - operation 27, 29
 - organization 29
 - replacement array 24
 - row definition 29
 - row deletion 31
- byte multiplexer channel
 - addresses 37
 - data rate 37
 - maximum number 37
 - subchannels 39
- byte-oriented operands 2, 15

- channel bus controller 44
- Channel Indirect Data Addressing feature 37
- channel logout 44, 46
- channel-not-operational error condition 45
- channel retry 44
- Channel-to-Channel Adapter 37
- channels
 - addressing 37
 - block multiplexer 37
 - byte multiplexer 37
 - Channel-to-Channel Adapter 37
 - data rates 37
 - Extended Channels feature 37
 - general description 36
 - interference 38
 - I/O device attaching 40
 - maximum number 37
 - priorities 41
 - recovery features 45
 - selector 37
 - Two-Byte Interface feature 37
 - UCW assignment 39
- CLEAR CHANNEL instruction 45
- clock comparator 2
- command retry 44
- common segment facility 34
- comparison tables
 - Models 165, 165 II, and 168, and 3033 hardware features 131
 - OS/VS support 140
- comparison with Model 168 126
- compatibility 9
- components of the processor complex 6
- console 49
- control registers 11
- control storage 17
- cooling, processor 7
- CPU timer 2
- cycle time
 - buffer storage 27
 - control storage 17
 - local storage 17
 - processor 14
- DAT hardware (see dynamic address translation)
- degradation machine check 21
- directors 37
- dynamic address translation 32
 - addresses translated 32
 - instruction nullification 32
 - segment table origin address saving 34
 - translation lookaside buffer 32
- ECC checking 25
- execution function 16
- exigent machine check conditions 23
- extended channel logout 46
- Extended Channels feature 37
- extended control mode 11
- external damage machine check 20
- features
 - optional 12
 - standard 11
- fixed processor storage locations 18

- imprecise interruptions 14
- initial microprogram load 52, 53
- instruction nullification 32
- instruction preprocessing function 14
- instruction processing damage interruption 23
- instruction retry 17, 45
- instructions
 - overlap of preparation and execution 14
 - standard set 11
- interface inoperable channel error condition 45
- interleaving, processor storage 24
- internal performance 1, 4
- interruptions
 - imprecise 14
 - machine check 18
 - masking machine check 22
 - other than machine check 18
- interval timer 2
- INVALIDATE PAGE TABLE ENTRY instruction 34
- I/O devices attachable 3
- I/O operation retry 45
- IPL, system state after 23

- limited channel logout 44
- local storage 17
- low address protection 13

- machine check code 20
- machine check interruptions 18
- main storage (see processor storage)
- maintenance and retry function 46
- monitoring feature 2
- motor generator set 7
- multiprocessing
 - advantages
 - loosely coupled 96
 - tightly coupled 91
 - buffer intercommunication 103
 - channel set switching 83
 - channels 83
 - configuration for the 3033 80
 - definition 78
 - loosely coupled 79
 - tightly coupled 79
 - differences between uniprocessor and multiprocessor hardware 98
 - floating storage addressing 87
 - instructions 98
 - interprocessor hardware communication 103
 - interprocessor programmed communication 101
 - I/O devices 83
 - models 88
 - oscillator assignment 88
 - performance 95
 - planning considerations 105
 - power control 90
 - prefixing 99
 - processor
 - addressing 100
 - features 80
 - storage 82, 87
 - programming support 78, 79, 82
 - recording and diagnostic programs 104
 - SET PREFIX instruction 99
 - SIGNAL PROCESSOR instruction 101

- STORE CPU ADDRESS instruction 100
- STORE PREFIX instruction 99
- time-of-day clock 100
- valid system configurations 88
 - 3036 Console 85
 - 3038 Multiprocessor Communication Unit 84
- MVS-dependent instructions 13
- MVS/System Extensions program product 123
- optional features 12
- OS/VS2 support
 - MVS support 123
 - SVS support 123
- processor-dependent logout area 19
- processor elements 10
- processor-independent logout area 18
- processor storage
 - cycle times 26
 - ECC 25
 - interleaving 24
 - logical storage elements 24
 - priority 25
 - reconfiguration 26
 - sizes 24
- processor storage control function 27
- program event recording 2
- programming systems support
 - MVS/System Extensions 123
 - OS/VS1 123
 - OS/VS2 MVS 123
 - OS/VS2 SVS 123
 - VM/Basic System Extensions 124
 - VM/System Extensions 124
 - VM/370 123
- PURGE TLB instruction 34
- reconfiguration, processor storage 26
- reference and change recording 32
- reloadable control storage 17
- repressible machine check conditions 20
- selector channels 37
- space requirements 6
- standard features 11
- storage
 - buffer 27
 - control 17
 - interleaving 24
 - local 17
 - processor (main) 24
 - reconfiguration 26
 - virtual 32
- store and fetch protection 25
- STO-stack 34
- subchannels 38
- system damage machine check 23
- system recovery machine check 20
- System/370 Extended Facility 12
- technology 8
- TEST PROTECTION instruction 13
- timer damage machine check 20
- time-of-day clock 2
- timing facilities damage machine check 21

trace unit 46
 translation lookaside buffer 32
 Two-Byte Interface feature 37

UCWs 38

virtual storage 32
 VM/Basic System Extensions 124
 Virtual Machine Extended Facility Assist 14
 VM/System Extensions 124
 VM/370 support 123, 124

3037 Power and Coolant Distribution Unit 7

3036 Console for the 3033 Processor Complex
 addresses 51
 alter/display frame 61
 cabling 49
 channel attachment 51
 components 49
 concurrent maintenance functions 71, 74
 configuration frame 62
 console configuration 62
 console processor 52
 control panel 52
 counter 52
 CRT display and keyboard 56
 director configuration frame 64
 diskette drive 50, 57
 diskettes provided 58, 59
 display frames 59
 event index frame 66
 fault locating tests 76
 frame verbiage 58
 functions 51
 index frame 59
 indicator/logout frames 62
 log index frame 66
 logouts 58
 maintenance mode configurations 72
 microcode loading 54, 55
 microdiagnostics frame 68
 mode control frame 66
 modem 50
 modes of operation 68
 normal configuration 70
 operating station components 50
 operator station 49, 68, 70, 72
 operator frame 60
 ports 69
 power control frame 66
 powering 54
 program frame 61
 recovery facilities 75
 remote maintenance operations 71
 security key on CRT display 57, 60
 service record file mode 59
 service support station 49, 69, 70, 72
 serviceability features 76
 standard I/O interface 55
 system activity frame 64
 system diskette 58
 system status recording 52, 64
 TP link frame 66
 voltage monitoring 76

SECTION 90: OS/VIRTUAL STORAGE 1 FEATURES

If required, the OS/Virtual Storage 1 Features Supplement (GC20-1752) should be inserted here.

This page intentionally left blank

SECTION 100: OS/VIRTUAL STORAGE 2 SINGLE VIRTUAL STORAGE (SVS) FEATURES

If required, the OS/Virtual Storage 2 Single Virtual Storage (SVS) Features Supplement (GC20-1753) should be inserted here.

This page intentionally left blank

SECTION 110: VIRTUAL MACHINE FACILITY/370 FEATURES

If required, the Virtual Machine Facility/370 Features Supplement (GC20-1757) should be inserted here.

This page intentionally left blank

This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name and mailing address:

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

What is your occupation? _____

Number of latest Newsletter associated with this publication: _____

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape


Please Do Not Staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.



POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department 824
1133 Westchester Avenue
White Plains, New York 10604

Fold and tape

Please Do Not Staple

Fold and tape

A Guide to the IBM 3033 Processor Complex, Attached Processor Complex, and Multi-processor Complex of System/370 Systems Printed in U.S.A. GC20-1859-4



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601