

R E M O V E T	PUB. NUMBERS	PAGE NUMBERS	TABS & DIVIDERS	BINDER SIZE	What is changed; why is it changed; how should this print package be used?	CRAY-2 Computer System Functional Description Manual Change Package HR-2000-01
★	HR-2000 C	Pages 1-1 through 1-7 Page 2-3/2-4 Page 3-5/3-6 Pages 4-1/4-2 Pages A-11/ A-12 Pages B-3 through B-5			Remove and shred all of Section 1. Remove and shred pages 2-3/2-4. Remove and shred pages 3-5/3-6. Print on blue paper. Remove and shred pages 4-1/4-2. Remove and shred pages A-11/A-12 . Remove and shred pages B-3 through B-5.	
★	HR-2000-01	Pages 1-1 through 1-7			Insert new Section 1 (pages 1-1 through 1-7) from the change package into the manual.	
★		Pages 2-3/2-4			Insert new Section 2 pages 2-3/2-4 from the change package into Section 2.	
★		Pages 3-5/3-6			Insert new Section 3 pages 3-5/3-6 from the change package into Section 3. Print on blue paper.	
★		Pages 4-1/4-2			Insert new Section 4 pages 4-1/4-2 from the change package into Section 4.	
★		Pages A-11/ A-12			Insert new Appendix A pages A-11/A-12 from the change package into Appendix A.	
★		Pages B-3 through B-5			Insert new Appendix B pages B-3 through B-5 from the change package into Appendix B.	

1. INTRODUCTION

The CRAY-2 computer system is a powerful, general-purpose computer system with extremely high processing rates. Scalar and vector capabilities in a multiprocessing environment combined with integrated foreground processing achieve these high rates.

1.1 CRAY-2 COMPUTER SYSTEM FEATURES

The CRAY-2 computer system mainframe contains either two or four independent Background Processors, each more powerful than a CRAY-1 computer system processor. Featuring a clock-cycle time faster than any other computer system available, each of these processors offers exceptional scalar and vector processing capabilities. The Background Processors can operate independently on separate jobs or concurrently on a single problem. The very high speed Local Memory integral to each Background Processor is available for temporary storage of vector and scalar data.

Common Memory is one of the most important features of the CRAY-2 computer system. It consists of 256 or 512 Mwords in dynamic memory, or 64 or 128 Mwords in static memory, 64-bits long, randomly accessible from any of the Background Processors and from any of the data channels. The memory is arranged in quadrants with either 64 or 128 interleaved banks. All memory access is performed automatically by the hardware. Any user may use all or part of the memory not being used by the operating system.

Control of network access equipment and the high-speed disk drives is integral to the CRAY-2 computer system mainframe hardware. A single Foreground Processor coordinates the data flow between the system's Common Memory and all the external devices across either two or four high-speed I/O channels. The synchronous operation of the Foreground Processor with the Background Processors and the external devices provides a significant increase in data throughput.

The most important CRAY-2 computer system features are:

- Extremely large directly addressable Common Memory
- Fastest cycle time available in a computer system
- Scalar, vector, and multiprocessing combined in one system
- Integral Foreground Processor

- Elegant architecture
- Extremely high reliability
- High density memory chips and extremely fast silicon logic chips
- Liquid immersion cooling

1.1.1 PHYSICAL CHARACTERISTICS


The CRAY-2 computer system mainframe is elegant in appearance as well as in architecture (see figure 1-1). The memory, computer logic, and DC power supplies are integrated into a compact mainframe composed of 14 vertical columns arranged in a 300° arc.

The upper part of each column contains a stack of logic modules and the lower part contains power supplies for the system. Total cabinet height, including the power supplies, is 45 in. (114.3 cm); the diameter of the mainframe is 53 in. (134.6 cm). Thus, the "footprint" of the mainframe is a mere 16 ft² (1.49 m²).

An inert fluorocarbon liquid circulates in the mainframe cabinet in direct contact with the integrated circuit packages. This liquid immersion cooling technology allows for the small size of the CRAY-2 computer system mainframe and is thus largely responsible for the high computation rates.

Significant CRAY-2 computer system physical characteristics are:

- Occupies only 16 ft² (1.49 m²) of floor space
- Stands 45 in. (114.3) high, diameter is 53 in. (134.6 cm)
- Contains 14 columns arranged in a 300° arc
- Contains 3-dimensional modules
- Contains liquid immersion cooling
- Contains chilled water heat exchange



1353

Figure 1-1. CRAY-2 Computer System Mainframe

1.1.2 ARCHITECTURE AND DESIGN

In addition to the cooling technology, the extremely high processing rates are achieved by a balanced integration of scalar and vector capabilities and a large Common Memory in a multiprocessing environment.

Significant architectural components of the CRAY-2 computer system include the following:

- Two or four independent Background Processors capable of vector and scalar operation. Synchronization of the Background Processors is achieved through the Foreground Processor and semaphore flags in the Background Processors.
- 256 or 512 Mwords of dynamic Common Memory, or 64 or 128 Mwords of static Common Memory
- A foreground system that controls and monitors system operation, including:
 - A Foreground Processor for system supervision
 - Two or four high-speed synchronous communication channels
 - Up to 40 I/O devices
 - Disk controllers to control up to 36 disk storage units (DSUs)
 - Two or four Common Memory ports for data transfer
 - Two or four Background Processor ports to allow Foreground Processor control
 - External I/O controllers (from one to as many as four per channel)
 - HSX controllers (two maximum per channel)

The identical Background Processors each contain registers and functional units to perform both vector and scalar operations. The single Foreground Processor supervises the Background Processors. The large Common Memory complements the processors and provides architectural balance, thus assuring extremely high throughput rates (see figure 1-2).

Shown in figure 1-2 is the four-processor model. The two-processor versions have two high-speed synchronous communication channels. The contents of a channel are the same in each version of the system.

On-site maintenance is possible through the maintenance control console.

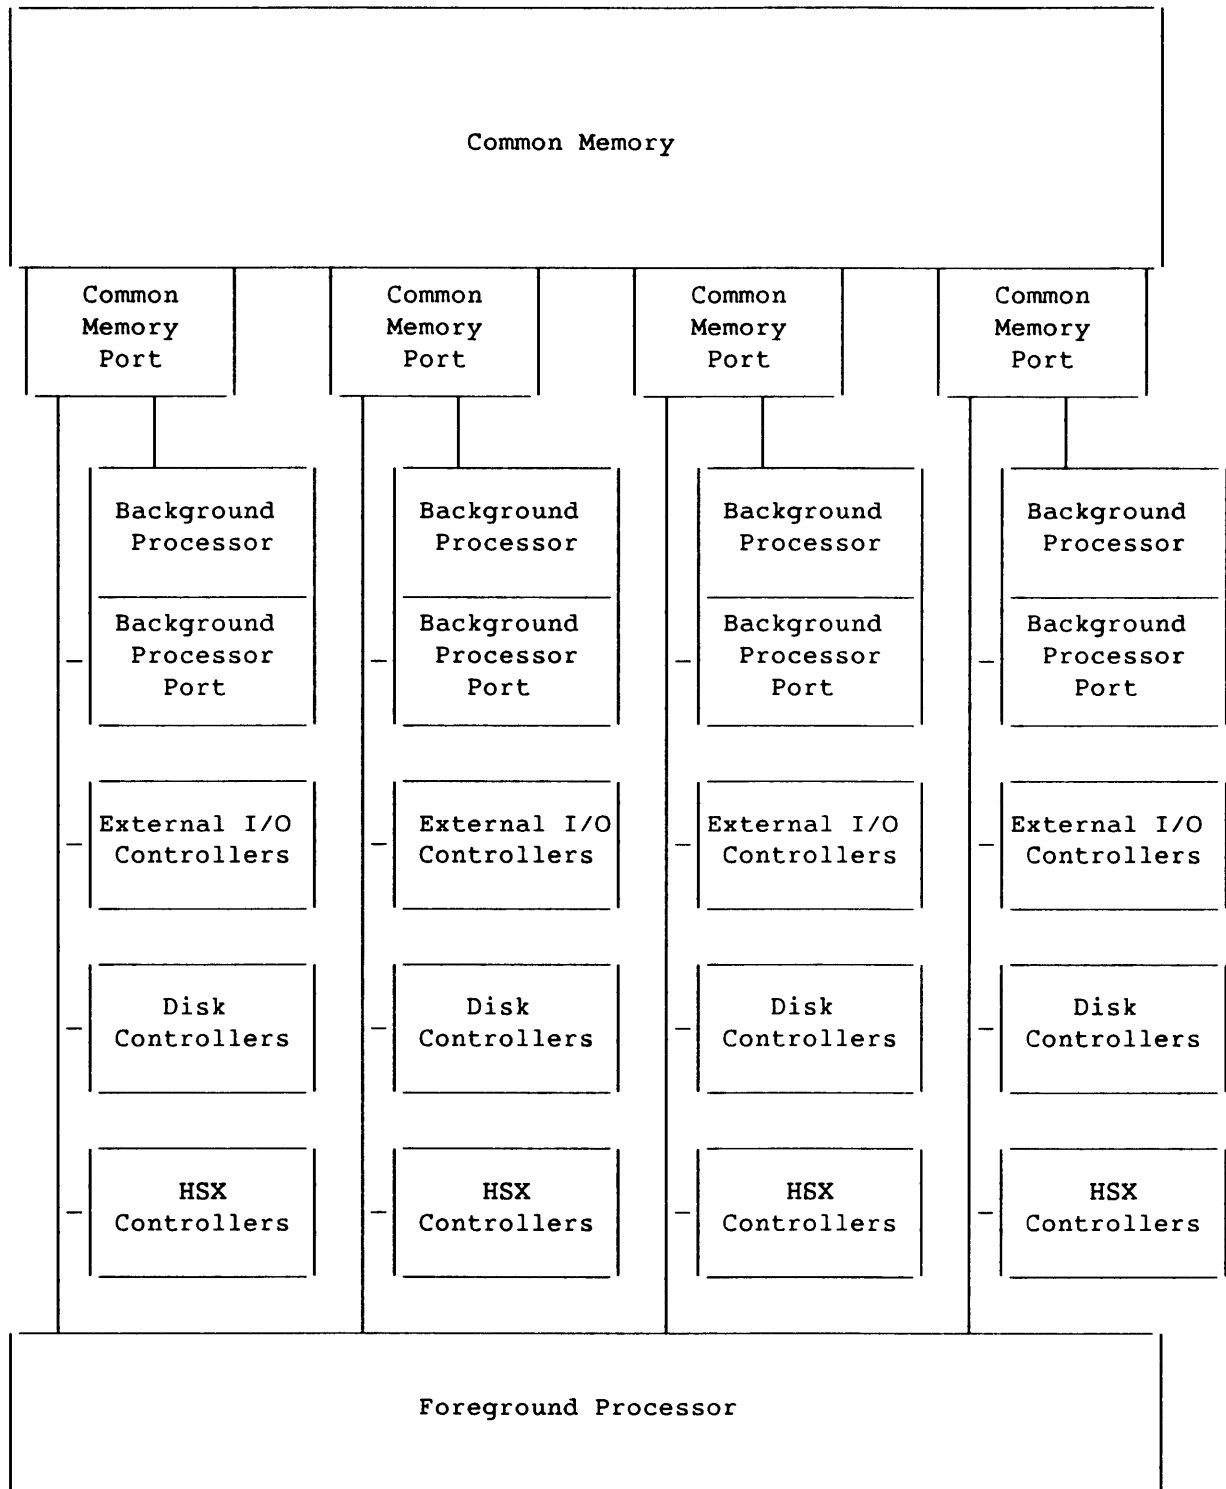


Figure 1-2. CRAY-2 Four Background Processor Computer System Mainframe Configuration

1.2 CONVENTIONS

This manual uses the following conventions:

<u>Convention</u>	<u>Description</u>
<i>lowercase italics</i>	Variable information
X or x or x	An ignored value
<i>n</i>	An unknown variable value
(Xx)	The contents of a register designated by the Xx value
Register bit designators	Numbered right to left as powers of 2, starting with 2 ⁰ .

Unless otherwise indicated, numbers in this manual are decimal numbers. Octal numbers are indicated with an 8 subscript. Exceptions are instruction parcels in instruction buffers and instruction forms which are given in octal without the subscript.

1.2.1 EXAMPLES

Illustrations of the above conventions.

<u>Example</u>	<u>Description</u>
Transmit (Ak) to Si	Transmit the contents of the A register specified by the k designator to the S register specified by the i designator
167ixk	Machine instruction 167 where the j register designator is not used and is an ignored value
Read n words from memory	Read an unknown variable number of words from memory. You can read, within the stated restrictions, as few or many words from memory as you wish.
Bit 2 ⁶³ of an S or V register	Value represents the most significant bit

<u>Example</u>	<u>Description</u>
Bit 2 ³¹ of an A register	Value represents the most significant bit
VM register element	The VM register contains 64 bits, each corresponding to a word element in a Vector register. Bit 2 ⁶³ corresponds to element 0, bit 2 ⁰ corresponds to element 63.

1.3 ORGANIZATION

This manual is organized into the following sections:

<u>Section</u>	<u>Description</u>
1	Contains the introduction to this manual
2	Describes the CRAY-2 computer system Background Processor. The registers, functional units, and algorithms used are described.
3	Provides detailed information on the CAL instructions that operate on the CRAY-2 computer system. Each machine instruction can be represented symbolically in Cray Assembly Language (CAL) Version 2. The instructions are listed octally in a box format that provides the Cray Assembly Language (CAL) Version 2 syntax format, an operand if required, a brief description of each instruction, and the machine instruction. Following the boxed information is a detailed description of the instruction and an example using the instruction.
4	Describes the CRAY-2 Common Memory, phased memory access, and single-error correction/double-error detection (SECCED)
5	Describes the CRAY-2 Foreground System, which handles the I/O
Appendix A	Lists the symbolic machine instructions by function. The octal machine code can be used as an index when referring to section 3 for a detailed description of the instruction.
Appendix B	Contains the CRAY-2 system configuration specification sheets

Instruction buffers

Each Background Processor has a buffer with eight independent fields to allow program loops to execute without additional Common Memory references. Programs can loop within the instruction buffer using any of the branch instructions.

Each independent field contains 16 or 32 words. The total instruction buffer size is 128 or 256 words.

The next sequential instruction out of the instruction buffer or a branch out of the instruction buffer discards the oldest data field and replaces it with 16 or 32 words of new data.

Instruction issue

Background instructions are translated in several steps and are allowed to issue sequentially by an instruction issue control mechanism. The words are disassembled into 16-bit parcels that are placed in a queue where the translation occurs. The instruction issue process involves checking the reservation flags for the registers and functional unit involved in the instruction sequence. The parcel waits in issue position in the instruction queue until all required resources are free.

Instruction parcels and 16-bit constants are intermixed in the instruction queue. The constant parcels are passed through the instruction queue without test.

2.1.2 REAL-TIME CLOCK

Each Background Processor has a 64-bit register that counts continuously at the clock period rate. This count value determines the passage of real time to an accuracy of 1 clock period (CP). The real-time clocks in the Background Processors are synchronized at deadstart. Instruction 115 reads the real-time clock.

2.1.3 SEMAPHORE FLAGS

To synchronize Common Memory references, eight semaphore flags in the background system interlock Common Memory references when multiple Background Processors are executing a single job. One semaphore flag is assigned to each currently active job in the background system. A Background Processor assigned to a job is assigned a semaphore flag at the same time.

The Background Processor uses four instructions in synchronizing its Common Memory references: 004, 005, 006, and 007. A 004 or 005 instruction requests the semaphore flag when the Background Processor program is accessing a Common Memory area that can interfere with other processors assigned to the job. The branch instruction results determine when the processor has exclusive access to this Common Memory area. The program must clear the semaphore flag to release the Common Memory area to another processor assigned to the same job.

2.1.4 COMMON MEMORY FIELD PROTECTION

At execution time each object program has a designated field of Common Memory holding instructions and data. The foreground functions specify the field limits when the object program is loaded and initiated. Field limits are contained in the Base Address (BA) register and the Limit Address (LA) register.

All memory addresses contained in the object program code are relative to the base address beginning the defined field. An object program cannot read or alter any Common Memory location with an absolute address lower than the base address. Each object program reference to Common Memory is checked against the limit and base addresses to determine if the address is within the assigned bounds.

Base Address register

Each Background Processor has a 32-bit BA register. The BA register defines the lower boundary of the Common Memory address field. The Foreground Processor enters data into this register while the Background Processor is in idle mode. The data remains in the register for the duration of the Background Processor computation period.

Each Common Memory reference from the Background Processor includes the addition of the BA register contents to the other parts of the memory reference base address. All Background Processor references to Common Memory are relative to the base address boundary.

Limit Address register

Each Background Processor has a 32-bit LA register. The LA register defines the upper boundary of the Common Memory address field. The Foreground Processor enters data into this register while the Background Processor is in idle mode. The data remains in this register for the duration of the Background Processor computation period.

INSTRUCTIONS 000 - 001

Result	Operand	Description	Machine Instruction
err		Error exit	000x00
exit		Normal exit	000x01
exit	<i>exp</i>	Normal exit	000x <i>jk</i>
CMR		Hold issue on memory busy	001xxx

Instructions 000 and 001 stop the current program sequence, place the Background Processor in idle mode, and set the Exit Mode and Idle Mode flags in the Background Port Status register. The 6-bit *jk* value is entered into the Background Port Status register.

Examples:

Code Generated	Location	Result	Operand	Comment
	1	10	20	35
000000		err		
000001		exit		
000004		exit	4	

INSTRUCTION 002

Result	Operand	Description	Machine Instruction
r, a_j	a_k	Register jump to (a_k) with return address to a_j	002ixk
j	a_k	Register jump to (a_k), value in a_k erased	002kxk

Instruction 002 stops the current program sequence and begins a new sequence at a computed parcel address read from the A_k register. The parcel address for the next instruction in the current program sequence is entered into the A_j register.

Examples:

Code Generated	Location	Result	Operand	Comment
	1	10	20	35
002102		r, a1	a2	
002101		j	a1	

4. COMMON MEMORY

Common Memory contains 256 or 512 Mwords of dynamic memory, or 64 or 128 Mwords of static memory. The memory consists of either 64 or 128 banks. Each 72-bit word consists of 64-data bits and 8 error-correction bits.

Common Memory is organized into quadrants with 32 banks in each quadrant. The 64 Mword version has 16 banks per quadrant. Each memory quadrant has a data path to each of the Common Memory ports. A Background Processor and a foreground communication channel are connected to each Common Memory port. The total memory bandwidth of a four-processor system is 64 Gbits/s. The total memory capacity is now equal to 17 Gbits.

The Foreground Processor, Background Processors, external I/O devices, and disk controllers share Common Memory. Common Memory contains program code for the Background Processors, data for problem solution, and Foreground Processor system tables.

4.1 MEMORY ADDRESSING

A word in memory is addressed by 32 bits. The low-order 2 bits select the quadrants and the next 5 bits select the bank. The 64-Mword system uses 4 bits for bank select. Figure 4-1 shows the format of the memory address for Common Memory.




Figure 4-1. Memory Address for Common Memory

4.2 MEMORY ACCESS

The Background Processors are locked into a phased access time scheme with the memory quadrants through the Common Memory ports. Through its Common Memory port, a Background Processor can access any given quadrant but only in the processor's own phase time, that is, every fourth clock period (CP). If a Background Processor requests a quadrant out of its phase time, the request is delayed until the correct time.

For example, assume the Background Processors are A through D, and the quadrants are 0 through 3. Also assume processor A is locked into quadrant 0 at phase time 0. If processor A references quadrant 0 at phase time 1, it must wait until the next phase time 0 (CP 4) to have access to memory in that quadrant.

Memory banks in a quadrant share a data path to each Common Memory port. Because of the phased access time between the quadrants and the Common Memory ports, however, only 1 bank accesses the path in a given 4-CP time slot. Because 2 banks never compete for the same data path in the same time slot, each bank functionally has an independent path to each of the four Common Memory ports.

4.3 MEMORY CONFLICTS

To prevent memory conflicts, each memory bank in the dynamic system has two Bank Busy flags. Each bank is divided logically into two or four pseudobanks. This enables quicker access to the half of the bank that is not busy. When a bank has been accessed it sets both of its busy flags. A long count busy applies to the pseudo bank that is actually busy, while a short count busy applies to the pseudo bank that is not. If the bank is busy, the quadrant sends a rejected signal to the requesting memory port. The requesting port retries the data.

The static memory, being much faster, does not require the pseudo bank arrangement. One bank busy is used per bank.

4.4 MEMORY BACKUP

Memory back-up occurs when too many memory references arrive at a single memory quadrant. Each Common Memory port has four quadrant buffers, one for each quadrant. Each buffer can hold two memory references for its memory quadrant. Therefore, references can continue to the memory port when the reference is not in the proper phase time. When a quadrant buffer in a memory port is filled, and another reference to that quadrant is made, the memory port begins a back-up procedure.

The memory port back-up procedure stops instruction issue for the associated Background Processor if that processor is making a memory reference. Vector streams initiated in the Background Processor and associated with a Common Memory reference are held.

After all references have been submitted for retry, stop issue is released allowing additional references to issue. A conflict during the retry process causes the back-up procedure to begin again at the point the conflict occurred; which could be the original back-up reference or another reference buffered during backup.

Common Memory writes

Result	Operand	Description	Machine Instruction
(<i>exp</i>)	s_i	Writes (s_i) to Common Memory at location exp	067ixx $m_1 m_2$
(a_k)	s_i	Writes (s_i) to Common Memory at location (a_k)	063ixk
(a_k, exp)	s_i	Writes (s_i) to Common Memory at location (a_k)+ exp	065ixk $m_1 m_2$
(a_j, a_k)	s_i	Writes (s_i) to Common Memory at location (a_j)+(a_k)	061ijk
(a_j, a_k)	v_i	Writes (v_i) to Common Memory location (a_j) incremented by (a_k)	071ijk
(a_k, v_j)	v_i	Scatters (v_i) to Common Memory locations (a_k)+(v_j)	073ijk

A.7.2 LOADS

Several instructions can be used to load data from memory into registers.

Local Memory reads

Result	Operand	Description	Machine Instruction
a_i	[exp]	Reads from location exp in Local Memory to a_i	044ixx m_l
a_i	[a_k]	Reads from location to a_k in Local Memory to a_i	046ixk
s_i	[exp]	Reads from location exp in Local Memory to s_i	054ixx m_l
s_i	[a_k]	Reads from location to a_k in Local Memory to s_i	056ixk
v_i	[a_k]	Reads from Local Memory location (a_k) to v_i	074ixk

Complete Memory references

Result	Operand	Description	Machine Instruction
CMR		Hold issue on memory busy	001xxx

CPU Features	
Number of CPUs	4
Clock speed	4.1 ns
Common memory size	512 Mwords or 256 Mwords
Common memory chip type	Dynamic MOS
Number of quadrants	4
Number of banks	128
Number of common memory ports	4
Number of foreground channels	4
Maximum number of I/O devices	40
Maximum number of disk storage devices	36
Maximum number of HSX controllers	8
Maximum number of external I/O controllers	16

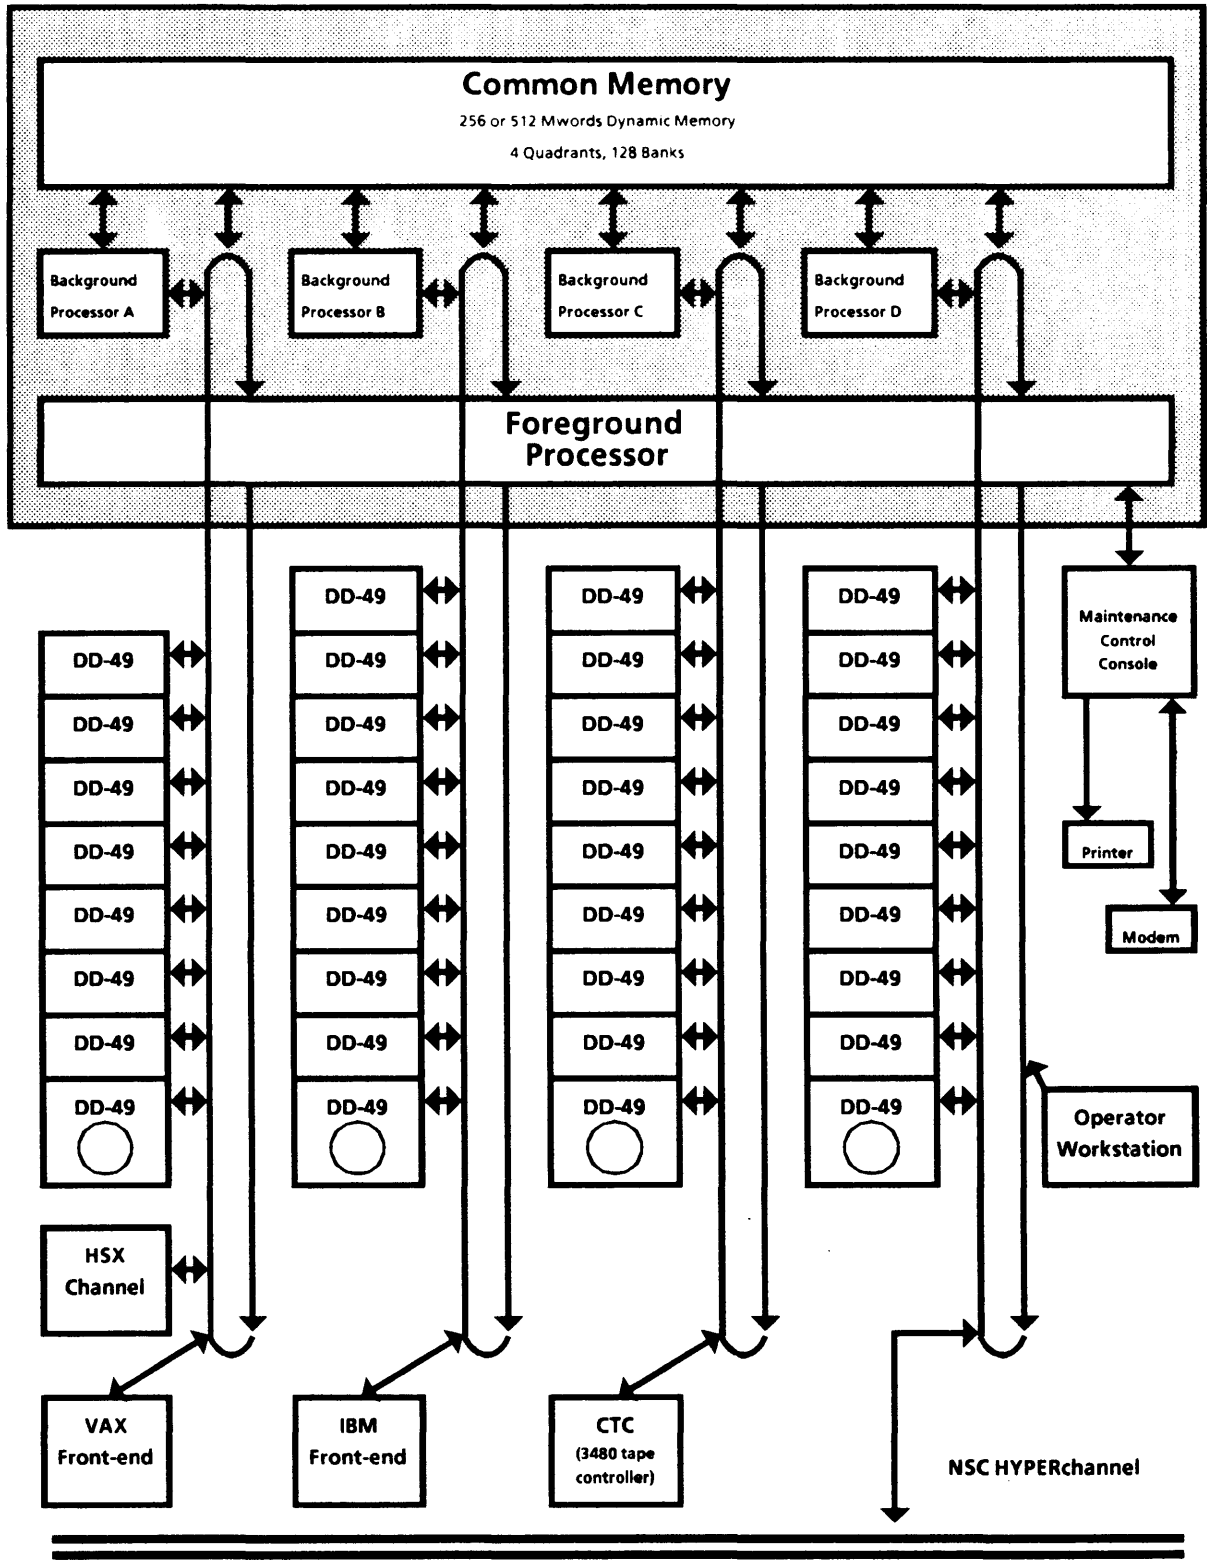
CPU Features continued	
Number of columns	14
Arc	300°
Floor space	16 ft ² (1.49 m ²)
Weight	5500 lb (2495 kg)

Functional Units (register units)
<small>Available per Background Processor</small>
Address functional units: <ul style="list-style-type: none"> ● Add/subtract (A) ● Multiply (A)
Scalar functional units: <ul style="list-style-type: none"> ● Integer <ul style="list-style-type: none"> -Add/subtract (S) -Population/parity (S) -Leading zero count (S) ● Shift (S) ● Logical (S)
Vector functional units: <ul style="list-style-type: none"> ● Integer <ul style="list-style-type: none"> -Add/subtract (S) -Shift (S) -Population/parity (S) -Leading zero count (S) -Compressed iota (S and V) ● Logical (S and V)
Floating-point functional units: <ul style="list-style-type: none"> ● Add/subtract (S and V) ● Multiply, reciprocal, and square root (S and V)

Register Type Available per Background Processor	Quantity	Size
Address (A)	8	32 bits
Scalar (S)	8	64 bits
Vector (V)	8	64 elements (64 bits per element)
Local Memory (used for register save)	1	16K 64-bit words

Support Equipment Required per CRAY-2 Computer System	Number of Units Needed
Reservoir	1
M-pod	1
S-pod	1
Motor-generator sets	3
Maintenance Control Console	1

CRAY-2/4-256 or 4-512 System Block Diagram





Copyright© 1987 by CRAY RESEARCH, INC.