

CRAY® C90 Series Functional Description Manual

HR-04028-0A

Cray Research, Inc.

Copyright © 1993 by Cray Research, Inc. This manual or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Research, Inc.

Autotasking, CF77, CRAY, Cray Ada, CRAY Y-MP, CRAY-1, HSX, MPGS, SSD, SUPERSERVER, UniChem, UNICOS, and X-MP EA are federally registered trademarks and CCI, CFT, CFT2, CFT77, COS, CRAY APP, CRAY S-MP, CRAY X-MP, CRAY XMS, CRAY-2, Cray C++ Compiling System, Cray/REELlibrarian, CRInform, CRI/TurboKiva, CSIM, CVT, Delivering the power . . ., Docview, EMDS, IOS, OLNET, RQS, SEGLDR, SMARTE, SUPERCLUSTER, SUPERLINK, and Trusted UNICOS are trademarks of Cray Research, Inc.

AEGIS and Apollo are trademarks of Apollo Computer Inc. Amdahl is a trademark of Amdahl Corporation. AOS is a trademark of Data General Corporation. Apollo and Domain are trademarks of Apollo Computer Inc. CDC is a trademark of Control Data Corporation. DEC, DECnet, PDP, VAX, VAXcluster, and VMS are trademarks of Digital Equipment Corporation. Delta Series is a trademark of Motorola, Inc. ECLIPSE is a trademark of Data General Corporation. E-Systems is a trademark of E-Systems, Inc. Ethernet is a trademark of Xerox Corporation. Fluorinert Liquid is a trademark of 3M. Honeywell is a trademark of Honeywell, Inc. HYPERchannel and NSC are trademarks of Network Systems Corporation. IBM is a trademark of International Business Machines Corporation. LANlord is a trademark of Computer Network Technology Corporation. NetBlazer and Telebit are trademarks of Telebit Corporation. Siemens is a trademark of Siemens Aktiengesellschaft of Berlin and Munich, Germany. SPARC and SPARCstation are a trademarks of SPARC International, Inc. OpenWindows, Sun-3, Sun-4, and SunOS are trademarks of Sun Microsystems, Inc. TRAK is a trademark of TRAK Microwave Corporation. UltraNet is a trademark of Ultra Network Technologies, Inc. Unisys is a trademark of Unisys Corporation. UNIX is a trademark of UNIX System Laboratories, Inc. The UNICOS operating system is derived from the UNIX System Laboratories, Inc. UNIX System V operating system. UNICOS is also based in part on the Fourth Berkeley Software Distribution under license from The Regents of the University of California. X Window System is a trademark of the Massachusetts Institute of Technology.

Requests for copies of Cray Research, Inc. publications should be directed to:

CRAY RESEARCH, INC.
Distribution
2360 Pilot Knob Road
Mendota Heights, MN 55120
800-284-2729 extension 35907

Comments about this publication should be directed to:

CRAY RESEARCH, INC.
Hardware Publications and Training
890 Industrial Blvd.
Chippewa Falls, WI 54729

Record of Revision

Each time this manual is updated with a change packet, a change to part of a text page is indicated by a change bar in the margin directly opposite the change. A change bar in the footer of a text page indicates that most, if not all, of the text is new. A change bar in the footer of a page composed primarily of a table and/or figure may indicate that a change was made to that table/figure or, it could indicate that the entire table/figure is new. Change packets are assigned a numerical designator, which is indicated in the publication number on each page of the change packet.

Each time this manual is fully revised and reprinted, all change packets to the previous version are incorporated into the new version, and the new version is assigned an alphabetical revision level, which is indicated in the publication number on each page of the manual. A revised manual does not usually contain change bars.

REVISION

DESCRIPTION

| | |
|----|--|
| | March 1992. Original printing. |
| 0A | April 1993. Technical changes and additions were made to include information about the CRAY C92A, CRAY C94A, CRAY C94, and CRAY C98 computer systems. Section 2 was updated to include information on the bit matrix multiply functional unit and new CRAY C90 series mainframe configurations. Section 3 was updated to include information on the DCA-3, TCA-2, HCA-5, and UTC-1 channel adapters. Information was added to Section 4 about the SSD-E/32i solid-state storage device. Information on the DD-62, RD-62, and disk array disk product was added to Section 5. |

PREFACE

The *CRAY C90 Series Functional Description Manual* provides information about the basic functions, design, and architecture of the CRAY C92A, CRAY C94A, CRAY C94, CRAY C98, and CRAY C916 computer systems and their associated peripheral devices. This manual is written primarily for Cray Research, Inc. (CRI) customers and people who desire a basic overview of the system.

This manual is divided into the following tabbed sections:

Section 1, "System Overview," introduces and describes the CRAY C90 series system components and support equipment.

Section 2, "Mainframe," describes the basic hardware architecture and CPU instructions of the mainframe. Specification sheets are included at the end of this section.

Section 3, "I/O Subsystem," describes the basic architecture and functions of the input/output subsystem (IOS). A specification sheet is included at the end of this section.

Section 4, "SSD Solid-state Storage Devices," describes the basic architecture of the SSD solid-state storage device model E (SSD-E) and the SSD-E/32i solid-state storage device. Specification sheets are included at the end of this section.

Section 5, "Peripheral Equipment," describes the function of the disk drives and network interface equipment used by the CRAY C90 series computer systems. Specification sheets are included at the end of this section.

Section 6, "Software Overview," provides an overview of the software available for the CRAY C90 series computer systems.

For the readers' convenience, a glossary is included. It defines many of the common abbreviations used and terminology associated with CRAY C90 series computer systems.

The following conventions are used throughout this manual.

| <u>Convention</u> | <u>Description</u> |
|--------------------------|---|
| Lowercase <i>italic</i> | Variable information. |
| X or x or <i>x</i> | An unused value. |
| n | A specified value. |
| (value) | The contents of the register or memory location designated by value. |
| Register bit designators | Register bits are numbered from right to left as powers of 2. Bit 2^0 corresponds to the least significant bit of the register. One exception is the vector mask register. The vector mask register bits correspond to a word element in a vector register; bit 2^{63} corresponds to element 0 and bit 2^0 corresponds to element 63. Another exception is when the contents of the 32 1-bit semaphore registers are loaded into an S register. SM0 goes into S register bit position 2^{63} , SM1 goes into S register bit position 2^{62} , and so on. |
| Number base | All numbers used in this manual are decimal, unless otherwise indicated. Octal numbers are indicated with an 8 subscript. Exceptions are register numbers, the instruction parcel in instruction buffers, and instruction forms, which are given in octal without the subscript. |

The following list provides examples of the preceding conventions.

| <u>Example</u> | <u>Description</u> |
|-------------------------------------|--|
| Transmit (<i>Ak</i>) to <i>Si</i> | Transmit the contents of the A register specified by the <i>k</i> field to the S register specified by the <i>i</i> field. |
| 167 <i>ixk</i> | Machine instruction 167. The <i>x</i> indicates that the <i>j</i> field is not used. |
| Read n words from memory | Read a specified number of words from memory. |
| Bit 2^{63} | The value represents the most significant bit of an S register or element of a V register. |
| 1000 ₈ | The number base is octal. |

CONTENTS

1 SYSTEM OVERVIEW

| | |
|--|------|
| Mainframe | 1-2 |
| I/O Subsystem | 1-4 |
| SSD Solid-state Storage Devices | 1-5 |
| Disk Storage Units | 1-5 |
| Tape Drives and Controllers | 1-6 |
| Operator and Maintenance Workstations | 1-6 |
| MWS-E Functions | 1-7 |
| OWS-E Functions | 1-8 |
| Network Interfaces | 1-8 |
| Power and Cooling Support Equipment | 1-9 |
| CRAY C92A and CRAY C94A Computer Systems | 1-9 |
| Power Equipment | 1-9 |
| Cooling Equipment | 1-10 |
| Warning and Control System | 1-11 |
| CRAY C94 and CRAY C98 Computer Systems | 1-12 |
| Power Equipment | 1-12 |
| Cooling Equipment | 1-13 |
| Warning and Control System | 1-14 |
| CRAY C916 Computer Systems | 1-15 |
| Power Equipment | 1-15 |
| Cooling Equipment | 1-15 |
| Warning and Control System | 1-18 |
| System Configurations | 1-18 |
| CRAY C92A Computer System Configurations | 1-19 |
| CRAY C94A Computer System Configurations | 1-21 |
| CRAY C94 Computer System Configurations | 1-23 |
| CRAY C98 Computer System Configurations | 1-25 |
| CRAY C916 Computer System Configurations | 1-27 |

2 MAINFRAME (continued)

| | |
|--|------|
| CPU Shared Resources | 2-1 |
| Central Memory | 2-1 |
| I/O Section | 2-3 |
| Interprocessor Communication Section | 2-3 |
| Real-time Clock | 2-4 |
| CPU Computation Section | 2-4 |
| Operating Registers | 2-5 |
| Address Registers | 2-6 |
| Scalar Registers | 2-6 |
| Vector Registers | 2-7 |
| Functional Units | 2-7 |
| Address Functional Units | 2-8 |
| Scalar Functional Units | 2-8 |
| Vector Functional Units | 2-9 |
| Floating-point Functional Units | 2-11 |
| Functional Unit Operations | 2-13 |
| Logical Operations | 2-13 |
| Integer Arithmetic | 2-14 |
| Floating-point Arithmetic | 2-16 |
| CPU Control Section | 2-30 |
| Exchange Mechanism | 2-30 |
| Exchange Sequence | 2-30 |
| Exchange Package | 2-31 |
| Instruction Fetch Sequence | 2-38 |
| Instruction Issue | 2-38 |
| Programmable Clock | 2-38 |
| Status Registers | 2-38 |
| Performance Monitor | 2-38 |
| Parallel Processing Features | 2-39 |
| Pipelining and Segmentation | 2-39 |
| Functional Unit Independence | 2-42 |
| Vector Processing | 2-42 |
| Advantages of Vector Processing | 2-43 |
| V Register Functions | 2-43 |

2 MAINFRAME (continued)

| | |
|--|------|
| Vector Instructions | 2-44 |
| Vector Chaining | 2-44 |
| Multiprocessing and Multitasking | 2-46 |
| Autotasking | 2-47 |
| CPU Instructions | 2-49 |
| Notational Conventions | 2-49 |
| Instruction Formats | 2-50 |
| 1-parcel Instruction Format with Discrete j and k Fields | 2-50 |
| 1-parcel Instruction Format with Combined j and k Fields | 2-51 |
| 2-parcel Instruction Format with Combined i, j, k, and m Fields | 2-52 |
| 3-parcel Instruction Format with Combined m and n Fields | 2-52 |
| Special Register Values | 2-53 |
| Special CAL Syntax Forms | 2-54 |
| Monitor Mode Instructions | 2-54 |
| Program Range | 2-55 |
| CPU Instruction Summary | 2-55 |
| Functional Units Instruction Summary | 2-56 |
| Functional Instruction Summary | 2-56 |
| Register Entry Instructions | 2-57 |
| Transfers into A Registers | 2-57 |
| Transfers into S Registers | 2-57 |
| Transfers into V Registers | 2-58 |
| Transfers into Semaphore Registers | 2-59 |
| Interregister Transfer Instructions | 2-59 |
| Transfers to A Registers | 2-59 |
| Transfers to S Registers | 2-60 |
| Transfers to V Registers | 2-61 |
| Transfers to Intermediate Registers | 2-61 |
| Transfers to Shared Registers | 2-61 |
| Transfers to Status Registers | 2-62 |
| Transfer to Vector Mask Register | 2-62 |
| Transfer to Vector Length Register | 2-62 |

2 MAINFRAME (continued)

| | |
|---|------|
| Memory Transfer Instructions | 2-63 |
| Bidirectional Memory Transfers | 2-63 |
| Memory References | 2-63 |
| Writes | 2-64 |
| Reads | 2-64 |
| Integer Arithmetic Instructions | 2-65 |
| 32-bit Integer Arithmetic | 2-66 |
| 64-bit Integer Arithmetic | 2-66 |
| Bit Matrix Multiply | 2-67 |
| Floating-point Arithmetic Instructions | 2-67 |
| Floating-point Range Errors | 2-67 |
| Floating-point Addition and Subtraction | 2-68 |
| Floating-point Multiplication | 2-68 |
| Reciprocal Approximation | 2-69 |
| Logical Operation Instructions | 2-69 |
| Logical Products | 2-70 |
| Logical Sums | 2-70 |
| Logical Exclusive ORs | 2-71 |
| Logical Equivalence | 2-71 |
| Vector Mask | 2-72 |
| Merge | 2-72 |
| Shift Instructions | 2-73 |
| Bit Count Instructions | 2-74 |
| Scalar Population Count | 2-74 |
| Vector Population Count | 2-74 |
| Population Parity Count | 2-75 |
| Scalar Leading Zero Count | 2-75 |
| Vector Leading Zero Count | 2-75 |
| Branch Instructions | 2-75 |
| Unconditional Branch Instructions | 2-76 |
| Conditional Branch Instructions | 2-76 |
| Return Jump | 2-77 |
| Normal Exit | 2-77 |
| Error Exit | 2-77 |

2 MAINFRAME (continued)

| | |
|--|------|
| Monitor Mode Instructions | 2-78 |
| Channel Control | 2-78 |
| Set Exchange Address | 2-78 |
| Set Real-time Clock | 2-78 |
| Set Cluster Number | 2-79 |
| Programmable Clock Interrupt | 2-79 |
| Operand Range Error Interrupt | 2-79 |
| Interprocessor Interrupt | 2-80 |
| Breakpoint Interrupt | 2-80 |
| Performance Counters | 2-80 |
| CRAY C90 Series Mainframe Specifications | 2-81 |

3 I/O SUBSYSTEM

| | |
|---|------|
| I/O Cluster | 3-1 |
| I/O Processor | 3-2 |
| I/O Buffers | 3-3 |
| Low-speed and High-speed Channels | 3-3 |
| Channel Adapters | 3-4 |
| CCA-1 Channel Adapter | 3-4 |
| DCA-1 Channel Adapter | 3-5 |
| DCA-2 Channel Adapter | 3-5 |
| DCA-3 Channel Adapter | 3-5 |
| HCA-3 and HCA-4 Channel Adapters | 3-6 |
| HCA-5 Channel Adapter | 3-7 |
| TCA-1 Channel Adapter | 3-7 |
| TCA-2 Channel Adapter | 3-7 |
| UTC-1 Channel Adapter | 3-8 |
| Workstation Interfaces | 3-8 |
| Programmable Real-time Interrupt | 3-9 |
| IOS Model E Specifications | 3-11 |

4 SSD SOLID-STATE STORAGE DEVICES

| | |
|-------------------------------------|-----|
| SSD-E | 4-1 |
| Physical Description | 4-1 |
| Memory | 4-2 |
| Mainframe Data Transfers | 4-3 |
| IOS-E Data Transfers | 4-3 |
| MWS-E and SSD-E Transfers | 4-4 |
| SSD-E/32i | 4-4 |
| Physical Description | 4-4 |
| Memory | 4-5 |
| Mainframe Data Transfers | 4-5 |
| IOS-E Data Transfers | 4-6 |
| MWS-E and SSD-E/32i Transfers | 4-6 |
| SSD Model E Specifications | 4-7 |
| SSD-E/32i Specifications | 4-9 |

5 PERIPHERAL EQUIPMENT

| | |
|--|------|
| Disk Controller Units and Disk Storage Units | 5-1 |
| Disk Drives | 5-1 |
| DD-60 Disk Drive | 5-1 |
| DD-61 Disk Drive | 5-6 |
| DD-62 Disk Drive | 5-7 |
| RD-62 Disk Drive | 5-8 |
| Disk Array | 5-8 |
| DS-41 Disk Subsystem | 5-9 |
| DS-40 Disk Subsystem | 5-11 |
| DD-49 Disk Drive | 5-13 |
| Network Interfaces | 5-14 |
| FEI-1 Front-end Interface | 5-14 |
| Fiber-optic Link | 5-14 |
| FEI-3 Front-end Interface | 5-15 |
| Direct Network Connections | 5-16 |
| High Performance Parallel Interface (HIPPI) | 5-16 |
| DEC VAX Supercomputer Gateway | 5-17 |

5 PERIPHERAL EQUIPMENT (continued)

| | |
|---|------|
| DD-60 Disk Drive Specifications | 5-19 |
| DD-61 Disk Drive Specifications | 5-21 |
| DD-62 Disk Drive Specifications | 5-23 |
| RD-62 Disk Drive Specifications | 5-25 |
| DA-60 Disk Drive Specifications | 5-27 |
| DA-62 Disk Drive Specifications | 5-29 |
| DS-40 and DS-40D Disk Subsystem Specifications | 5-31 |
| DS-41, DS-41D, and DS-41R Disk Subsystem Specifications ... | 5-33 |
| DD-49 Disk Drive Specifications | 5-35 |
| Front-end Interface Specifications | 5-37 |
| FOL-3 Fiber-optic Link Specifications | 5-39 |

6 SOFTWARE OVERVIEW

| | |
|-------------------------------|------|
| UNICOS Operating System | 6-1 |
| Multiprocessing | 6-2 |
| Macrotasking Feature | 6-2 |
| Microtasking Feature | 6-2 |
| Autotasking Feature | 6-3 |
| CF77 Compiling System | 6-3 |
| C Compiler | 6-4 |
| Pascal | 6-5 |
| Cray Assembler | 6-5 |
| Cray Ada Environment | 6-5 |
| Cray Allegro CL | 6-6 |
| Subroutine Libraries | 6-6 |
| Utilities | 6-6 |
| Communications Software | 6-7 |
| Applications | 6-8 |
| Software Publications | 6-9 |
| UNICOS Operating System | 6-9 |
| Fortran | 6-9 |
| C | 6-9 |
| Pascal | 6-10 |
| Libraries | 6-10 |

6 SOFTWARE OVERVIEW (continued)

| | |
|-------------------------------|------|
| Utilities | 6-10 |
| Communications Software | 6-10 |
| Applications | 6-11 |
| Software Training | 6-11 |

FIGURES

| | |
|---|------|
| Figure 1-1. CRAY C90 Series Computer System Block Diagram | 1-3 |
| Figure 1-2. MWS-E and OWS-E Workstation Chassis | 1-6 |
| Figure 1-3. CRAY C92A and CRAY 94A Cooling System Configuration | 1-11 |
| Figure 1-4. CRAY C98 and CRAY C94 Cooling System Configuration | 1-14 |
| Figure 1-5. CRAY C916 Cooling System Configuration | 1-17 |
| Figure 2-1. CRAY C90 Series Mainframe Block Diagram ... | 2-2 |
| Figure 2-2. Integer Data Formats | 2-14 |
| Figure 2-3. 24-bit Integer Multiply Performed in a Floating-point Multiply Functional Unit | 2-15 |
| Figure 2-4. 32-bit Integer Multiply Performed in a Floating-point Multiply Functional Unit | 2-16 |
| Figure 2-5. Floating-point Data Format | 2-16 |
| Figure 2-6. Biased and Unbiased Exponent Ranges | 2-17 |
| Figure 2-7. Internal Representation of a Floating-point Number | 2-17 |
| Figure 2-8. Floating-point Add and Floating-point Multiply Range Errors | 2-19 |
| Figure 2-9. Floating-point Reciprocal Approximation Range Errors | 2-20 |
| Figure 2-10. Newton's Method for Approximating Roots | 2-24 |
| Figure 2-11. Vector Storage of a Bit Matrix | 2-27 |
| Figure 2-12. Matrix A and Matrix B | 2-28 |
| Figure 2-13. Matrix B and Matrix B ^t | 2-28 |
| Figure 2-14. Matrix C | 2-29 |
| Figure 2-15. Scalar Segmentation and Pipelining Example | 2-40 |

FIGURES (continued)

| | | |
|--------------|--|------|
| Figure 2-16. | Vector Segmentation and Pipelining Example | 2-41 |
| Figure 2-17. | Vector Chaining Example | 2-45 |
| Figure 2-18. | Vector Mask Bits | 2-49 |
| Figure 2-19. | General Instruction Format | 2-50 |
| Figure 2-20. | 1-parcel Instruction Format with Discrete j and k Fields | 2-51 |
| Figure 2-21. | 1-parcel Instruction Format with Combined j and k Fields | 2-51 |
| Figure 2-22. | 2-parcel Instruction Format with Combined i, j, k, and m Fields | 2-52 |
| Figure 2-23. | 3-parcel Instruction Format with Combined m and n Fields | 2-53 |
| Figure 3-1. | Cluster 0 Channel Connections | 3-2 |
| Figure 3-2. | Programmable Real-time Interrupt Signal Paths | 3-9 |
| Figure 5-1. | DD-60 Single-port Configurations | 5-3 |
| Figure 5-2. | DD-60 Daisy Chain Configuration | 5-4 |
| Figure 5-3. | DD-60 Alternate-path Configurations | 5-5 |
| Figure 5-4. | Disk Array Overview Block Diagram | 5-8 |

TABLES

| | | |
|------------|---|------|
| Table 1-1. | CRAY C92A Computer System Configurations | 1-20 |
| Table 1-2. | CRAY C94A Computer System Configurations | 1-22 |
| Table 1-3. | CRAY C94 Computer System Configurations . . . | 1-24 |
| Table 1-4. | CRAY C98 Computer System Configurations . . . | 1-26 |
| Table 1-5. | CRAY C916 Computer System Configurations | 1-28 |
| Table 2-1. | CRAY C90 Series Interrupt Modes | 2-33 |
| Table 2-2. | CRAY C90 Series Interrupt Flags | 2-35 |
| Table 2-3. | CRAY C90 Series Status Field Bit Assignments . . | 2-36 |
| Table 2-4. | CRAY C90 Series Operating Modes | 2-36 |

TABLES (continued)

| | | |
|------------|-------------------------------|------|
| Table 2-5. | Special Register Values | 2-54 |
| Table 4-1. | SSD-E Memory Sizes | 4-2 |

GLOSSARY

| | |
|----------------|-------|
| Glossary | Glo-1 |
|----------------|-------|

BIBLIOGRAPHY

| | |
|--------------------|-------|
| Bibliography | Bib-1 |
|--------------------|-------|

INDEX

| | |
|-------------|-------|
| Index | Ind-1 |
|-------------|-------|

1 SYSTEM OVERVIEW

The CRAY C90 series consists of five product lines: the CRAY C92A, CRAY C94A, CRAY C94, CRAY C98, and CRAY C916 computer systems. The naming convention for the CRAY C90 series is CRAY C9*nA*/*xy*, where *n*, *x*, and *y* represent the following numbers:

- *n* = maximum number of central processing units (CPUs) the mainframe can contain
- *x* = number of CPUs actually contained in the mainframe
- *y* = number of Mwords of central memory in the mainframe
- A = air cooled

The CRAY C90 series computer systems are powerful, general-purpose supercomputers. The large memory, fast clock speed, and powerful input/output (I/O) capabilities enable fast throughput, resulting in efficient use of supercomputing power. The CRAY C90 series computer systems achieve extremely high multiprocessing rates by efficiently using the scalar and vector processing capabilities of the multiple central processing units (CPUs), combined with the system's random-access memory (RAM) and shared registers.

With up to 16 powerful CPUs and up to 8 Gbytes (1 Gword) of central memory, each CRAY C90 series system is designed as a cost-effective solution for users with memory-constrained workloads.

A standard CRAY C90 series computer system consists of the following components:

- A mainframe
- An input/output subsystem model E (IOS-E)
- An optional solid-state storage device (SSD)
- Mass storage devices such as disk and tape drives
- A maintenance workstation model E (MWS-E)
- An operator workstation model E (OWS-E)
- Network interfaces
- Power and cooling support equipment

The following subsections introduce the system components. Subsequent tabbed sections provide more detailed information on the mainframe, IOS-E, SSD, peripheral devices, and system software. To simplify references to both SSD devices, the term *SSD* is used throughout this section to refer to the SSD-E/32i and SSD solid-state storage device model E (SSD-E), unless stated otherwise. Refer to “System Configurations” later in this section for more information about specific models in the CRAY C90 series.

Figure 1-1 is a block diagram of a typical CRAY C90 series system.

Mainframe

A CRAY C90 series mainframe contains an I/O section, an interprocessor communication section, central memory, and a variable number of CPUs. All CPUs in multiprocessor systems share the I/O section, interprocessor communication section, and central memory.

The mainframe is designed to deliver optimum overall performance. Separate registers and functional units support both integer and floating-point computations in both vector and scalar processing modes.

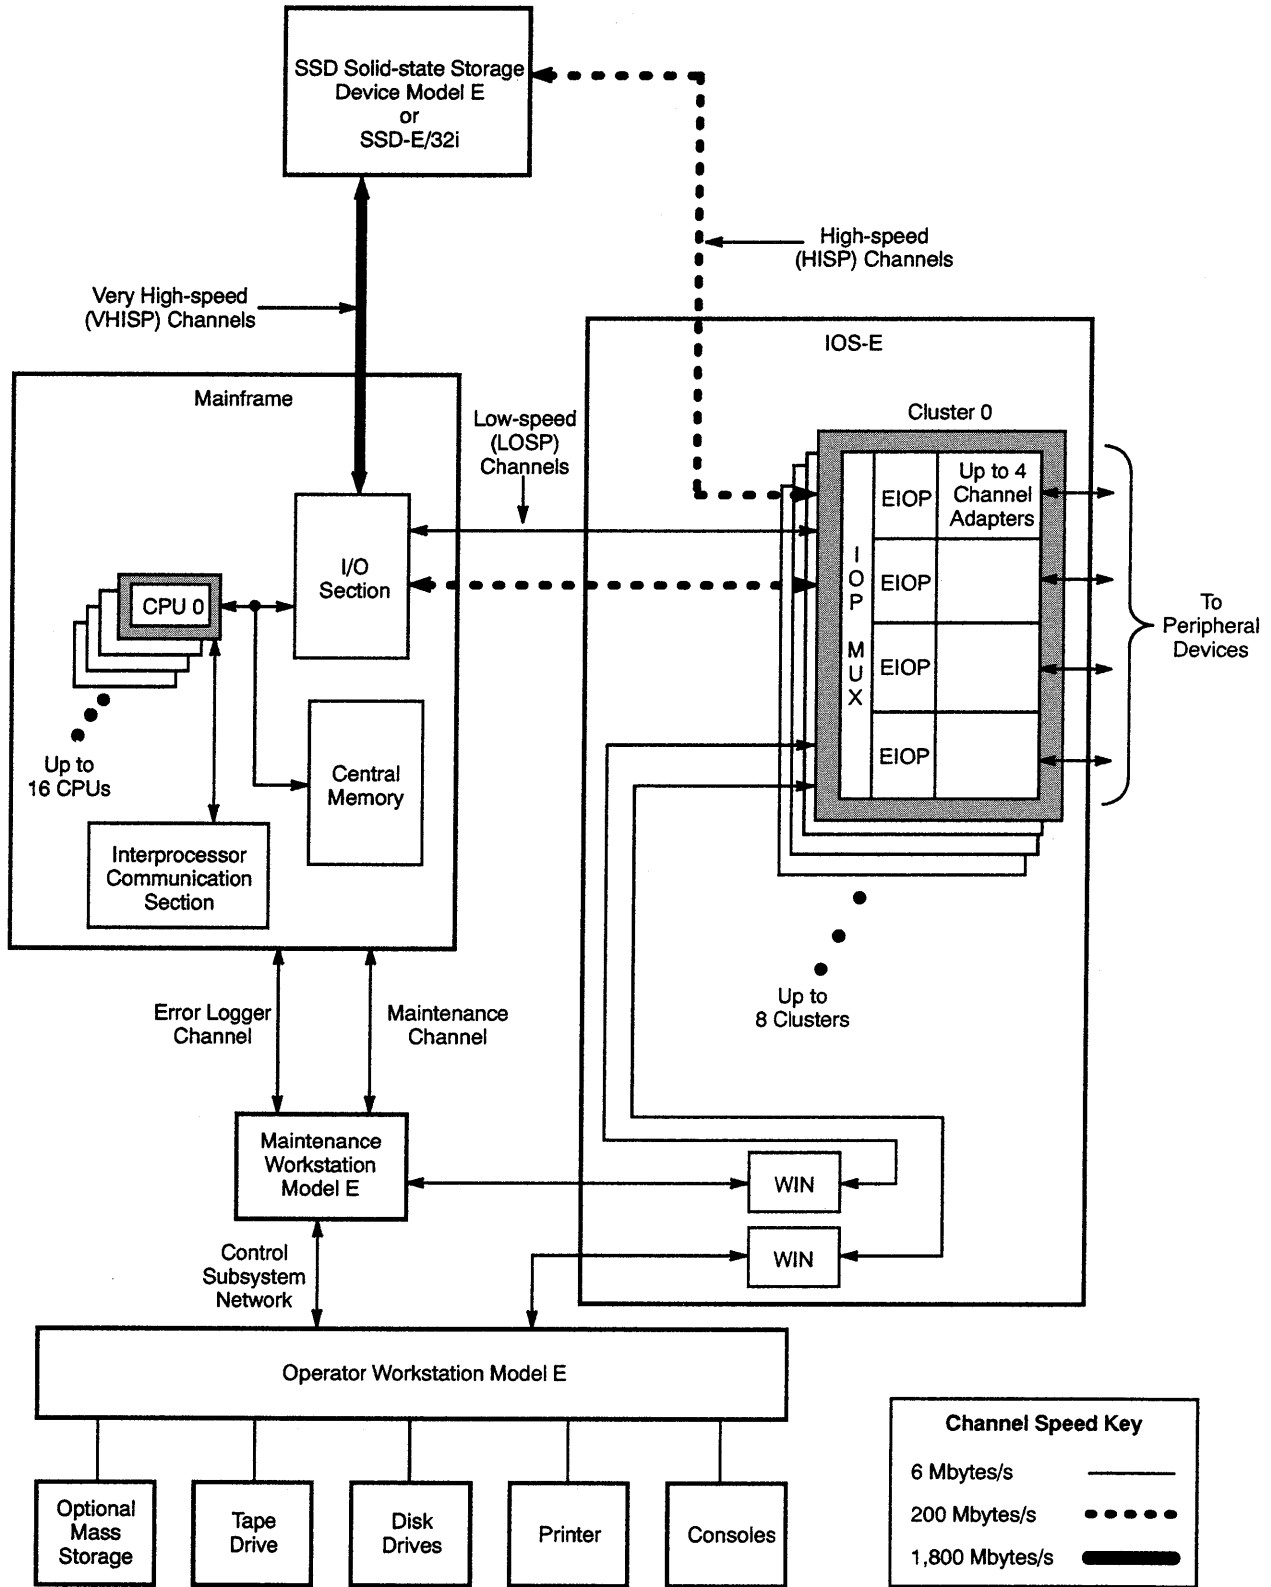
The I/O section provides high-speed data transfers to and from the IOS-E and SSD. The I/O section contains three types of I/O channels:

- 6-Mbyte/s low-speed (LOSP) channels
- 200-Mbyte/s high-speed (HISP) channels
- 1,800-Mbyte/s very high-speed (VHISP) channels

The LOSP channels carry control information between the mainframe and IOS-E. The HISP channels carry data between the mainframe and IOS-E and IOS-E and SSD. The VHISP channels carry data only between the mainframe and the SSD. The quantity of each channel type varies with different system configurations and depends on the quantity of CPUs and I/O clusters. Basically, each of the CPUs in the mainframe is configured with one LOSP channel and either two HISP channels or one VHISP channel.

The interprocessor communication section enables each mainframe CPU to synchronize operation and pass data with other CPUs. Central memory holds program code and data. Central memory is available in different sizes and configurations.

Each CPU has a control section and a computation section. The control section determines instruction issue and use of the computation section, central memory, and I/O resources. The computation section consists of operating registers and functional units.



A-11523

Figure 1-1. CRAY C90 Series Computer System Block Diagram

Vector processing uses a single instruction to perform multiple operations on sets of ordered data. Scalar processing is a sequential operation where one instruction produces one result. When two or more vector operations are chained together, two or more operations execute simultaneously. Therefore, the computational rate for vector processing greatly exceeds that of conventional scalar processing. Scalar operations complement the vector capability by providing solutions to problems not readily adaptable to vector techniques.

The start-up time for vector operations is short enough that vector processing is more efficient than scalar processing for vectors containing as few as two elements. This feature allows fast vector processing to be balanced with high-speed scalar processing.

Multiple-processor systems allow the use of multiprocessing or multitasking techniques. Multiprocessing allows several programs to be run concurrently on multiple CPUs of a single mainframe. Multitasking allows two or more parts of a program to run in parallel, sharing a common memory space.

Refer to Section 2, "Mainframe," for more information on the internal operation of the mainframe.

I/O Subsystem

All CRAY C90 series computer systems include an IOS-E. The IOS-E performs the following functions:

- Controls all data transfers between the mainframe or optional SSD and peripheral devices such as disk drives and communications networks.
- Buffers all data transfers between the mainframe or SSD and peripheral devices.
- Converts data to and from the formats used by peripheral devices.
- Detects and corrects certain types of data errors that occur during transfers.

The transfer rates between the IOS-E and other equipment vary. Data transfers between the IOS-E and either the mainframe or the SSD use HISP channels. The mainframe and IOS-E transfer control information using the 6-Mbyte/s channels. The transfer rate between the IOS-E and peripheral devices depends on the peripheral device.

The IOS-E consists of a variable number of I/O clusters and two workstation interfaces (WINs). Each I/O cluster contains one I/O processor multiplexer (IOP MUX), four auxiliary I/O processors (EIOPs), and up to 16 channel adapters. The IOP MUX controls data transfers between the IOS-E and mainframe or SSD. The four EIOPs control data transfers between the IOS-E and peripheral devices through the channel adapters. Each EIOP can support a maximum of four channel adapters. The number of clusters and channel adapters varies, depending on the number and types of peripheral devices in the system.

Workstation interfaces allow an OWS-E and MWS-E to control and monitor the operation of the I/O clusters. Refer to Section 3, "I/O Subsystem," for more information on the internal operation of the IOS-E.

SSD Solid-state Storage Devices

The SSD-E and the SSD-E/32i are optional high-performance devices used for temporary data storage. The SSD-E is contained in the same cabinet as the IOS-E. The SSD-E/32i consists of a single-coldplate module located in one of the cabinets containing the computer system. The SSD transfers data between the mainframe's central memory and the SSD through VHISP channels. The VHISP channel operates under mainframe program control. The SSD can also connect to the IOS-E by means of a HISP channel. The HISP channels operate under IOS-E program control. Refer to Section 4, "SSD Solid-state Storage Devices," for more information on the internal operation of the SSD-E and SSD-E/32i.

Disk Storage Units

The CRAY C90 series computer systems use Cray Research disk drives for mass data storage. A disk channel adapter (DCA-1, DCA-2, or DCA-3) provides an interface between the disk drives and an EIOP. The EIOP and the disk channel adapter can transfer data between the EIOP and multiple disk drives at full speed, even when all the drives are operating simultaneously. Refer to Section 3 for more information about channel adapters. Refer to Section 5 for more information about disk storage units.

Tape Drives and Controllers


The IOS-E provides an interface to tape drives and controllers. (Cray Research does not sell tape drives or controllers.) The TCA-1 channel adapter in the IOS-E connects to IBM compatible magnetic tape drives and controllers. Refer to Section 3, "I/O Subsystem," for more information on the channel adapters.

Operator and Maintenance Workstations

The MWS-E and OWS-E (Figure 1-2) are based on a Sun-4 370 SPARCstation, 12-slot chassis. The SPARC (Scalable Processor ARChitecture) is a SPARC International, Inc. version of the reduced instruction set computer (RISC) architecture. A VMEbus is provided in slots 4 through 12 of the workstations.

Both workstations run the SunOS 4.1.2 operating system and OpenWindows 3.0 software; the MWS-E also runs the MME maintenance diagnostic software release and the OWS-E also runs the OWS-E software release. The Sun operating system is an enhanced version of UNIX; it combines features of UNIX System Laboratories, Inc.'s System V UNIX and Berkely Software Distribution's version 4.3 UNIX. OpenWindows is a Sun system based on the OPEN LOOK standard and the X Window System.

The OWS-E is part of the Cray Research computer system. The MWS-E is owned by Cray Research; it enables Cray Research engineers to perform system maintenance independently of any customer activity on the Cray Research computer system.



A-10369

Figure 1-2. MWS-E and OWS-E Workstation Chassis

MWS-E Functions

The MWS-E provides an intelligent and dedicated platform for performing hardware maintenance, monitoring, and supporting of Cray Research computer systems.

The MWS-E is used to perform the following functions:

Offline diagnostic testing. Offline tests for the mainframe, IOS-E, SSD, and peripheral devices are loaded on the MWS-E. These diagnostics are used to verify proper hardware operation, to reproduce failures, and to isolate failures to the replaceable component.

Offline diagnostic listings. Listings are available online to assist Cray Research engineers in performing maintenance.

System deadstart and master clear.

IOS-E status. The MWS-E can read IOS-E status, read or write to local memory in an IOS-E processor (EIOP), and perform maintenance features such as deadstarting or master clearing an EIOP.

Hardware error logging. The error acquisition software program (EASE) records errors received through mainframe, IOS-E, and SSD error channels. EASE displays logged errors in an understandable format. The MWS-E also monitors system error channels to detect and log system errors such as double-bit memory errors.

Environmental monitoring. The MWS-E monitors the warning and control system (WACS) and responds to abnormal conditions. The WACS signals the Cray Research system to shut down for serious conditions and logs environmental variances that can later be used for failure analysis.

Remote support. The R3.0 Remote Support system provides a network connection to a remote location through a Telebit NetBlazer router and Microcom high-speed modem. The R3.0 release allows support personnel to dial into the site, log on the MWS-E, run maintenance tools, and monitor the Cray Research computer system.

Stand-alone disk testing (DD-40s, DD-41s, DD-60 series, and RDS-5s). The MWS-E serves as a stand-alone disk maintenance system for several disk drives sold with Cray Research computer systems. A disk drive supported in this manner can be removed from the system and serviced without the aid of system resources.

Stand-alone SSD testing. The data test channel that connects the MWS-E to the SSD enables you to test the SSD by using the low-speed (LOSP) data test channel. This channel enables you to test the SSD without using the high-speed (HISP) or very high-speed (VHISP) channels, which are dedicated to customer use and are normally connected to the mainframe and I/O subsystem.

SMARTE platform. The System Maintenance and Remote Testing Environment (SMARTE) is an online maintenance program used to perform hardware verification, error detection, error isolation, and automated degradation of faulty hardware components.

The OWS-E provides a dedicated workstation that Cray Research analysts and customer operators use to operate, administer, and monitor a Cray Research computer system. The OWS-E is also used for system boot, dump, clear, and troubleshooting operations and for software support and upgrades. For more information about the OWS-E, refer to the following publications:

- *OWS-E Operator Workstation Reference Manual*, publication number SG-3077
- *OWS-E Operator Workstation Operator's Guide*, publication number SG-3078
- *OWS-E Operator Workstation Administrator's Guide*, publication number SG-3079

Network Interfaces

The CRAY C90 series computer systems are designed to communicate easily and efficiently with front-end computer systems and computer networks.

Standard front-end interfaces (FEIs) connect the I/O channels of the IOS-E to front-end computer channels. These connections provide input data to the system and receive output from the system for distribution to peripheral equipment. An FEI compensates for differences in channel widths, machine word size, electrical logic levels, and control signals.

Some FEIs are housed in a stand-alone cabinet located near the host computer; others are installed directly into the front-end computer system. Operation of the FEI is transparent to both the front-end computer users and Cray Research system users.

An optional fiber-optic link (FOL-3 or FOL-4) is available for some FEIs to provide equipment separation distances of up to 13,120 ft (4,000 m). The FOL is installed between the IOS-E and FEI and provides complete electrical separation from the CRAY C90 series computer system.

Refer to “Network Interfaces” in Section 5 for more information on network interfaces.

Power and Cooling Support Equipment

The logic modules in the mainframe, IOS-E, and SSD require special equipment to supply electrical power and to remove heat. The following subsections define the power and cooling support equipment and the warning and control system (WACS) used with CRAY C90 series computer systems. Refer to Section 5, “Peripheral Equipment,” for power and cooling requirements for peripheral devices.

CRAY C92A and CRAY C94A Computer Systems

The CRAY C92A and CRAY C94A computer systems consist of one or two cabinets that house the mainframe, IOS-E, and SSD logic modules. The number of module cabinets in a CRAY C92A or CRAY C94A system can vary, depending on whether the system has an optional SSD-E. The 4200 (C92A) and 4400 (C94A) series cabinets always house the mainframe and IOS-E modules. If the system includes an optional SSD-E, an external cabinet houses the SSD-E logic modules. An external cabinet is not necessary for the SSD-E/32i solid-state storage device.

Power Equipment

This section describes the power distribution within the module cabinets used in CRAY C92A and C94A systems. The electrical differences between the various module cabinets are minimal.

The only electrical differences between the cabinets are the power supply and power bus configurations. The warning and control system (WACS) for each cabinet is also slightly different to account for the various power-supply configurations.

All power conditioning equipment for the logic and control circuitry is contained within the module cabinets. The module cabinets do not require motor-generator sets. Each module cabinet and each cooling unit contain a single drop cable that connects to standard commercial power. The following subsections describe the required input voltages and grounding systems.

The customer-supplied site power must include one of the following input voltages:

- 208-Vac, 60/50-Hz 3-phase
- 480-Vac, 60-Hz 3-phase
- 380-Vac, 50-Hz 3-phase
- 415-Vac, 50-Hz 3-phase

A wall circuit-breaker panel and power plug control the commercial power and feed it to the module cabinets and cooling units. The module cabinets each connect to a 200-A receptacle. The cooling units each connect to a 90-A receptacle.

The power supplies are contained within the mainframe chassis and IOS-E/SSD-E cabinet. The power supplies convert the voltages to the necessary DC voltages required for the logic modules.

Cooling Equipment

One or two cooling units, using room air or chilled water, cool a CRAY C92A or CRAY C94A system. One cooling unit is required for the mainframe cabinet, and one cooling unit is required for the IOS-E/SSD-E cabinet. The cooling unit is located in the computer room, approximately 2 ft (0.6m) from the cabinet it cools.

Figure 1-3 is a simplified diagram of the refrigeration system for a module cabinet. Cooling for each cabinet is accomplished by three systems: a dielectric-coolant system, a refrigerant system, and a chilled-water system or room air. The dielectric-coolant system contains a pump that circulates chilled dielectric fluid through each module. The dielectric fluid absorbs heat generated by the modules. The fluid then flows to a heat exchanger subassembly, where heat transfers from the dielectric fluid to the refrigerant system.

The refrigerant system contains a compressor that circulates the refrigerant. The refrigerant absorbs heat from the dielectric fluid and is then circulated through one of two condensers: one that is air cooled, and one that is water cooled. This dual-condenser design allows the system to be air cooled or water cooled without modification. The final stage of cooling transfers heat from the refrigerant system to room air or chilled water.

The internal isolation transformer and power supplies of the module cabinet are air cooled, regardless of whether the refrigerant system is air cooled or water cooled. Fans draw air in the front and sides of the cabinet. The air circulates around the isolation transformer and power supplies and exhausts out the top back of the cabinet.

The cooling unit transformer is also air cooled. Air enters the cooling unit through the back of the cabinet. Fans circulate air within the cooling unit cabinet and exhaust warm air out the top of the cabinet.

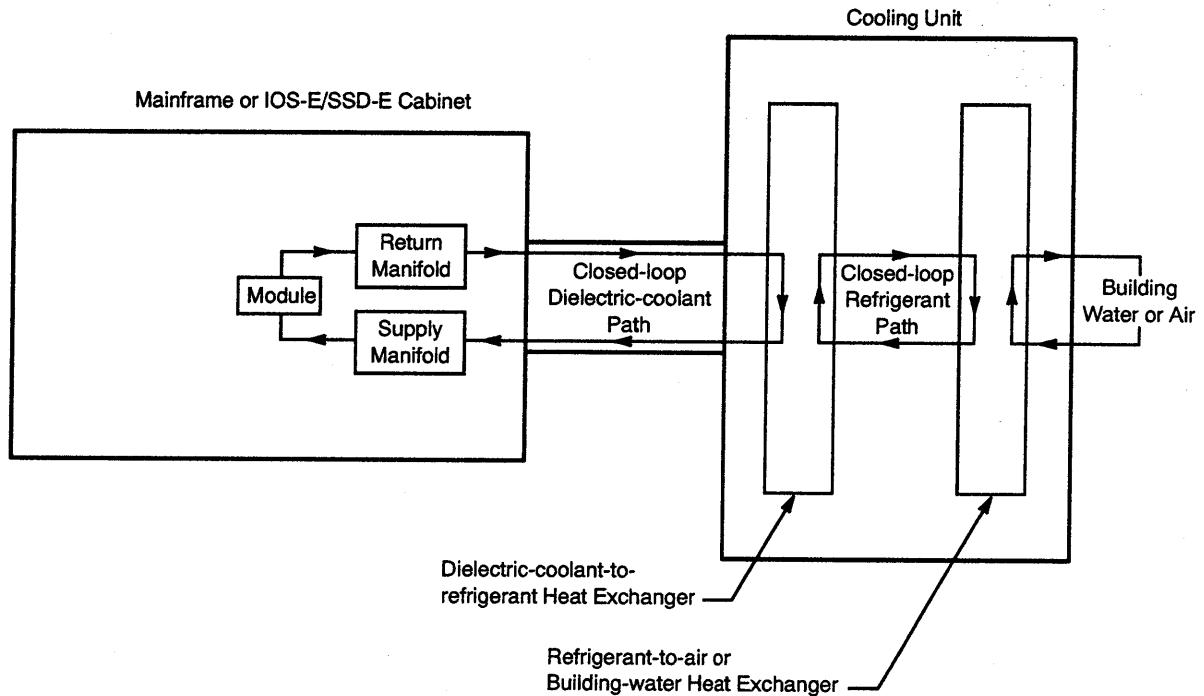


Figure 1-3. CRAY C92A and CRAY 94A Cooling System Configuration

Warning and Control System

This section describes the warning and control systems (WACS) that monitor and control refrigeration and power distribution for the various cabinets in the computer system. The WACS protects the equipment from damage by continuously monitoring environmental conditions such as temperature and pressure within the cabinet. The WACS can remove electrical power from the cabinet if warning or fault conditions exist. The WACS also reports the warning and fault conditions to a display window on the maintenance workstation model E (MWS-E).

The WACS consists of printed circuit boards and a system control panel. The WACS monitors the following conditions:

- Module temperature
- Power-supply voltage and current
- Dielectric-fluid pressure
- Inlet/outlet manifold temperatures
- AC input voltages
- Room humidity level using dewpoint monitors
- Internal voltages using self-tests
- Cooling unit dielectric-coolant level
- Smoke

The warning and control systems for the cabinets are almost identical. Some components (the power control board and display board) are slightly different to accommodate the various power-supply and busing configurations used in the different cabinets. The WACS operates on a 120- or 220-Vac, 60-Hz power source, or a 100- or 220-Vac, 50-Hz power source.

CRAY C94 and CRAY C98 Computer Systems

The CRAY C94 and CRAY C98 computer systems consist of one or two cabinets that house the mainframe, IOS-E, and SSD logic modules. The number of module cabinets in a CRAY C94 or CRAY C98 system can vary, depending on whether the system has an optional SSD-E. The 4600 (C94) and 4800 (C98) series cabinets always house the mainframe and IOS-E modules. If the system includes an optional SSD-E, an external cabinet houses the SSD-E logic modules. The SSD-E/32i solid-state storage device cannot be configured with the CRAY C94 or CRAY C98 computer systems.

Power Equipment

The motor-generator set (MGS-4) and power supplies provide electrical power to the logic modules and are housed in a stand-alone cabinet. The MGS is typically located in a separate power equipment room. The power supplies are housed in the mainframe cabinet.

An MGS uses power from commercial power to generate the proper voltage and frequency used by the mainframe power supplies. Customers must supply one of the following commercial power sources to the MGSs:

- 460 Vac, 3 phase, 60 Hz or
- 398 Vac, 3 phase, 50 Hz

The MGS supplies 208-Vac, 400-Hz power to the power supplies. The MGS also isolates the system from transients and fluctuations from commercial power.

The power supplies are housed in the mainframe chassis and IOS-E/SSD-E cabinet. The power supplies convert the 208-Vac, 400-Hz voltage to the necessary DC voltages required for the logic modules.

Cooling Equipment

The CRAY C94 and CRAY C98 computer systems use a heat exchanger unit (HEU) to transfer the heat energy from the dielectric-coolant that circulates through the logic modules to the refrigerant.

The CRAY C94 and CRAY C98 computer systems use refrigeration condensing unit RCU-9. The RCU dissipates the heat transferred from the HEU to customer-supplied chilled water.

NOTE: The RCU-9 and MGS-4 are configured with the majority of CRAY C90 series computer systems. However, some CRAY C90 series computer systems use different support equipment.

An HEU and RCU, along with customer-supplied chilled water, cool the computer system. The RCU is typically located in a separate equipment room.

Figure 1-4 is a simplified diagram of the cooling system for the CRAY C94 or CRAY C98 computer system. Cooling is accomplished by three systems: one or two closed-loop dielectric-coolant systems, one closed-loop refrigerant system, and a customer-supplied chilled water system.

Each closed-loop dielectric-coolant system contains a pump that circulates chilled dielectric coolant (such as Fluorinert Liquid) through each module and power-supply mounting plate. The dielectric coolant absorbs heat generated by the modules and power supplies. It then flows to a heat exchanger subassembly, where heat transfers from the dielectric coolant to the closed-loop refrigerant system.

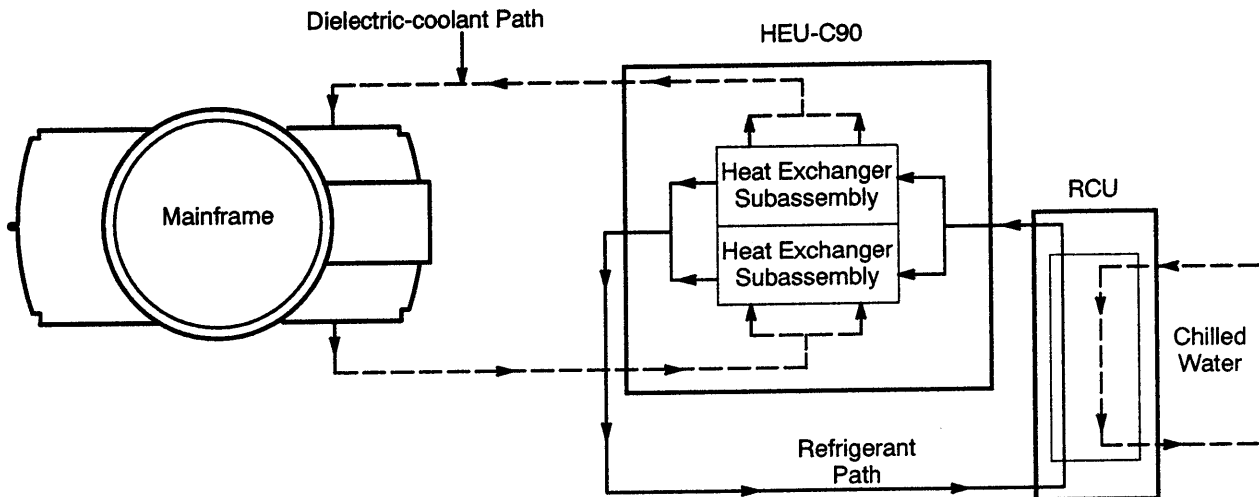


Figure 1-4. CRAY C98 and CRAY C94 Cooling System Configuration

Warning and Control System

The mainframe chassis and the IOS-E/SSD-E cabinet each contain a warning and control system (WACS). The WACS protects the equipment from damage by continuously monitoring environmental conditions such as temperature and pressure within the cabinet. The WACS can remove electrical power from the cabinet if warning or fault conditions exist. The WACS also reports the warning and fault conditions to the MWS-E.

The WACS consists of printed circuit boards and a system control panel. The WACS monitors the following conditions:

- Module temperature
- Power-supply voltage and current
- Dielectric-fluid pressure
- Inlet/outlet manifold temperatures
- AC input voltages
- Room humidity level using dewpoint monitors
- Internal voltages using self-tests
- HEU and RCU conditions
- Smoke

The WACS operates on a 120- or 220-Vac, 60-Hz power source, or a 100- or 220-Vac, 50-Hz power source. This power source is separate from the MGS power source.

CRAY C916 Computer Systems

A CRAY C916 computer system consists of two cabinets that house the mainframe, IOS-E, and SSD-E logic modules. The mainframe cabinet houses the CPU, memory, and clock logic modules. The IOS-E/SSD-E cabinet houses up to 8 clusters of IOS-E modules and the optional SSD-E modules. All cabinets use the same power and cooling technology as described in the following subsections.

Power Equipment

The motor-generator sets (MGSs) and power supplies provide electrical power to the logic modules and are housed in a stand-alone cabinet. The MGS is typically located in a separate power equipment room. The power supplies are housed in the mainframe cabinet.

An MGS uses commercial power to generate the proper voltage and frequency used by the mainframe power supplies. Customers must supply one of the following commercial power sources to the MGSs:

- 460 Vac, 3 phase, 60 Hz or
- 398 Vac, 3 phase, 50 Hz

The MGSs supply 208-Vac, 400-Hz power to the power supplies. The MGSs also isolate the system from transients and fluctuations from commercial power. The number of required MGSs varies among systems. Depending on the system configuration, a CRAY C916 computer system requires one or two MGSs. The CRAY C916 system uses two MGS-4s, an MGS-6, or an MGS-6A. MGS-4s are used in facilities where the installation of the MGS-6 or MGS-6A is not feasible or where MGS-4s have previously been installed. When two MGS-4s are used, a motor-generator parallel cabinet (MGPC) must be installed to combine the 400-Hz output of the two MGS-4s to produce a parallel frequency.

The power supplies are housed in the mainframe chassis and IOS-E/SSD-E cabinet. The power supplies convert the 208-Vac, 400-Hz voltage to the necessary DC voltages required for the logic modules. The number and types of power supplies are not optional for the customer.

Cooling Equipment

The CRAY C916 computer system uses two models of heat exchanger units (HEUs): HEU-C90 and HEU-E/S. The mainframe is connected to the HEU-C90, which has two pumps and two heat exchanger

subassemblies. The IOS-E/SSD-E cabinet is connected to the HEU-E/S, which contains a single pump and heat exchanger subassembly. The HEUs transfer heat between the dielectric coolant and the refrigerant.

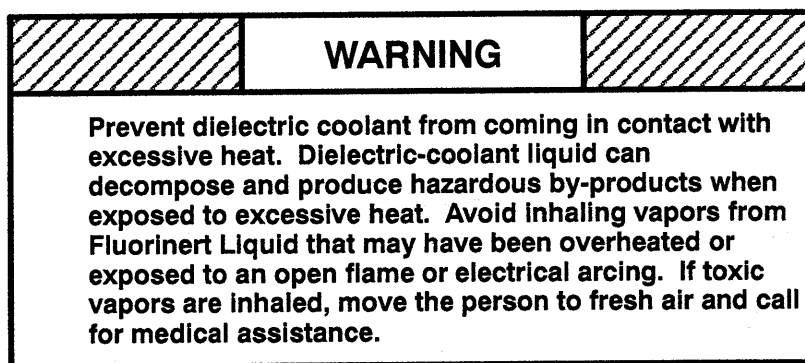
The CRAY C916 computer system uses the refrigeration condensing units RCU-5A or RCU-9. The RCU dissipates the heat transferred from the HEUs.

NOTE: The RCU-5A or RCU-9, and MGS-6 are configured with the majority of CRAY C916 computer systems. However, some CRAY C916 computer systems use different support equipment.

An HEU and RCU, along with customer-supplied chilled water, cool the CRAY C916 computer system. The RCU is typically located in a separate equipment power room.

Figure 1-5 is a simplified diagram of the cooling system for the CRAY C916 computer system. Cooling is accomplished by three systems: one or two closed-loop dielectric-coolant systems, one closed-loop refrigerant system, and a customer-supplied chilled water system.

Each closed-loop dielectric-coolant system contains a pump that circulates chilled dielectric coolant (such as Fluorinert Liquid) through each module and power-supply mounting plate. The dielectric coolant absorbs heat generated by the modules and power supplies. It then flows to a heat exchanger subassembly, where heat transfers from the dielectric coolant to the closed-loop refrigerant system.



The manual *Safe Use and Handling of Fluorinert Liquids*, Cray Research publication number HR-0306, provides specific guidelines and information regarding Fluorinert Liquid.




Figure 1-5. CRAY C916 Cooling System Configuration

The closed-loop refrigerant system contains a compressor that circulates the refrigerant. As previously mentioned, the refrigerant absorbs heat from the dielectric coolant. The refrigerant is then circulated through a condenser, where heat transfers to customer-supplied chilled water.

Cray Research recommends a water-supply temperature of approximately 50 °F (10 °C). Other chilled water specifications, such as flow rate and pressure-drop values, vary with different system configurations and actual water-supply temperatures. Cray Research provides additional information on these specifications during the site planning process.

Warning and Control System

The mainframe chassis and the IOS-E/SSD-E cabinet each contain a warning and control system (WACS). The WACS protects the equipment from damage by continuously monitoring environmental conditions such as temperature and pressure within the cabinet. The WACS can remove electrical power from the cabinet if warning or fault conditions exist. The WACS also reports the warning and fault conditions to the MWS-E.

The WACS consists of printed circuit boards and a system control panel. The WACS monitors the following conditions:

- Module temperature
- Power-supply voltage and current
- Dielectric-coolant pressure
- Dielectric-coolant flow rate
- Inlet/outlet manifold temperatures
- AC input voltages
- Room humidity level using dewpoint monitors
- Internal voltages using self-tests
- HEU and RCU conditions
- Smoke

The WACS operates on a 120- or 220-Vac, 60-Hz power source, or a 100- or 220-Vac, 50-Hz power source. This power source is separate from the MGS power source.

System Configurations

The various CRAY C90 series configurations accommodate a wide range of customer requirements and resources. Most models are field upgradeable. Customers can upgrade the number of CPUs, size of central memory, number or type of channel adapters, and so on.

The following specifications provide additional information about the various models of the CRAY C90 series product line.

CRAY C92A Computer System Configurations

Model Numbers

**CRAY C92A/164, CRAY C92A/1128,
CRAY C92A/264, CRAY C92A/2128**



CRAY
RESEARCH, INC.


Table 1-1. CRAY C92A Computer System Configurations

| Model Number | Mainframe Specifications | | | | | | | IOS Specifications | | SSD Memory Options in Mwords | |
|--------------|--------------------------|-----------------|--------------|-----------|--------------------------------|-------------|-----------|--------------------|--------------------|------------------------------|----------------------------|
| | Central Memory | | | | Maximum Number of I/O Channels | | | Number of CPUs | Number of Clusters | | Number of Channel Adapters |
| | Size in Mwords | No. of Sections | No. of Banks | Chip Size | 1,800 Mbyte/s | 200 Mbyte/s | 6 Mbyte/s | | | | |
| C92A/164 | 64 | 4 | 64 | 4 Mbit | 0 | 1 | 1 | 1 | 1 | 8 to 16 | N/A |
| C92A/1128 | 128 | 8 | 128 | 4 Mbit | 0 | 1 | 1 | 1 | 1 | 8 to 16 | N/A |
| C92A/264 | 64 | 4 | 64 | 4 Mbit | 1 | 2 | 2 | 2 | 1 to 2 | 8 to 32 | 32, 512, 1,024, or 2,048 |
| C92A/2128 | 128 | 8 | 128 | 4 Mbit | 1 | 2 | 2 | 2 | 1 to 2 | 8 to 32 | 32, 512, 1,024, or 2,048 |

CRAY C94A Computer System Configurations

Model Numbers

CRAY C94A/2128, CRAY C94A/4128



CRAY
RESEARCH, INC.


Table 1-2. CRAY C94A Computer System Configurations

| Model Number | Mainframe Specifications | | | | | | | IOS Specifications | | SSD Memory Options in Mwords | |
|--------------|--------------------------|-----------------|--------------|-----------|--------------------------------|-------------|-----------|--------------------|--------------------|------------------------------|----------------------------|
| | Central Memory | | | | Maximum Number of I/O Channels | | | Number of CPUs | Number of Clusters | | Number of Channel Adapters |
| | Size in Mwords | No. of Sections | No. of Banks | Chip Size | 1,800 Mbyte/s | 200 Mbyte/s | 6 Mbyte/s | | | | |
| C94A/2128 | 128 | 4 | 128 | 4 Mbit | 1 | 2 | 2 | 2 | 1 to 3 | 8 to 48 | 32, 512, 1,024, or 2,048 |
| C94A4128 | 128 | 4 | 128 | 4 Mbit | 2 | 4 | 4 | 4 | 1 to 3 | 8 to 48 | 32, 512, 1,024, or 2,048 |

CRAY C94 Computer System Configurations

Model Numbers

**CRAY C94/2128, CRAY C94/2256,
CRAY C94/4128, CRAY C94/4256**



CRAY
RESEARCH, INC.

Table 1-3. CRAY C94 Computer System Configurations

| Model Number | Mainframe Specifications | | | | | | | IOS Specifications | | SSD Memory Options in Mwords | |
|--------------|--------------------------|-----------------|--------------|-----------|--------------------------------|-------------|-----------|--------------------|--------------------|------------------------------|----------------------------|
| | Central Memory | | | | Maximum Number of I/O Channels | | | Number of CPUs | Number of Clusters | | Number of Channel Adapters |
| | Size in Mwords | No. of Sections | No. of Banks | Chip Size | 1,800 Mbyte/s | 200 Mbyte/s | 6 Mbyte/s | | | | |
| C94/2128 | 128 | 4 | 128 | 4 Mbit | 1 | 2 | 2 | 2 | 1 to 2 | 8 to 32 | 512, 1,024, or 2,048 |
| C94/2256 | 256 | 4 | 256 | 4 Mbit | 1 | 2 | 2 | 2 | 1 to 2 | 8 to 32 | 512, 1,024, or 2,048 |
| C94/4128 | 128 | 8 | 128 | 4 Mbit | 2 | 4 | 4 | 4 | 1 to 4 | 8 to 64 | 512, 1,024, or 2,048 |
| C94/4256 | 256 | 8 | 256 | 4 Mbit | 2 | 4 | 4 | 4 | 1 to 4 | 8 to 64 | 512, 1,024, or 2,048 |

CRAY C98 Computer System Configurations

Model Numbers

**CRAY C98/4256, CRAY C98/4512,
CRAY C98/8256, CRAY C98/8512**



CRAY
RESEARCH, INC.


Table 1-4. CRAY C98 Computer System Configurations

| Model Number | Mainframe Specifications | | | | | | | IOS Specifications | | SSD Memory Options in Mwords | |
|--------------|--------------------------|-----------------|--------------|-----------|--------------------------------|-------------|-----------|--------------------|--------------------|------------------------------|----------------------------|
| | Central Memory | | | | Maximum Number of I/O Channels | | | Number of CPUs | Number of Clusters | | Number of Channel Adapters |
| | Size in Mwords | No. of Sections | No. of Banks | Chip Size | 1,800 Mbyte/s | 200 Mbyte/s | 6 Mbyte/s | | | | |
| C98/4256 | 256 | 4 | 256 | 4 Mbit | 2 | 4 | 4 | 4 | 1 to 4 | 8 to 64 | 512, 1,024, or 2,048 |
| C98/4512 | 512 | 8 | 512 | 4 Mbit | 2 | 4 | 4 | 4 | 1 to 4 | 8 to 64 | 512, 1,024, or 2,048 |
| C98/8256 | 256 | 4 | 256 | 4 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 8 to 128 | 512, 1,024, or 2,048 |
| C98/8512 | 512 | 8 | 512 | 4 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 8 to 6128 | 512, 1,024, or 2,048 |

CRAY C916 Computer System Configurations

Model Numbers

**CRAY C916/8128, CRAY C916/8256, CRAY C916/8512, CRAY C916/81024,
CRAY C916/16128, CRAY C916/16256, CRAY C916/16512, CRAY C916/161024**



CRAY
RESEARCH, INC.

Table 1-5. CRAY C916 Computer System Configurations

| Model Number | Mainframe Specifications | | | | | | | IOS Specifications | | | SSD Memory Options in Mwords |
|--------------|--------------------------|-----------------|--------------|-----------|--------------------------------|-------------|-----------|--------------------|--------------------|----------------------------|------------------------------|
| | Central Memory | | | | Maximum Number of I/O Channels | | | Number of CPUs | Number of Clusters | Number of Channel Adapters | |
| | Size in Mwords | No. of Sections | No. of Banks | Chip Size | 1,800 Mbyte/s | 200 Mbyte/s | 6 Mbyte/s | | | | |
| C916/8128 | 128 | 4 | 512 | 1 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 15 to 128 | 512, 1,024, 2,048, or 4,096 |
| C916/8256 | 256 | 8 | 1,024 | 1 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 15 to 128 | 512, 1,024, 2,048, or 4,096 |
| C916/8512 | 512 | 4 | 512 | 4 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 15 to 128 | 512, 1,024, 2,048, or 4,096 |
| C916/81024 | 1,024 | 8 | 1,024 | 4 Mbit | 4 | 8 | 8 | 8 | 1 to 8 | 15 to 128 | 512, 1,024, 2,048, or 4,096 |
| C916/16128 | 128 | 4 | 512 | 1 Mbit | 4 | 16 | 16 | 16 | 1 to 16 | 15 to 256 | 512, 1,024, 2,048, or 4,096 |
| C916/16256 | 256 | 8 | 1,024 | 1 Mbit | 4 | 16 | 16 | 16 | 1 to 16 | 15 to 256 | 512, 1,024, 2,048, or 4,096 |
| C916/16512 | 512 | 4 | 512 | 4 Mbit | 4 | 16 | 16 | 16 | 1 to 16 | 15 to 256 | 512, 1,024, 2,048, or 4,096 |
| C916/161024 | 1,024 | 8 | 1,024 | 4 Mbit | 4 | 16 | 16 | 16 | 1 to 16 | 15 to 256 | 512, 1,024, 2,048, or 4,096 |

2 MAINFRAME

This section describes the major functional areas and special features of a CRAY C90 series mainframe and provides a summary of the Cray Assembly Language (CAL) instruction set. A CRAY C90 series mainframe specification sheet is included at the end of this section.

CPU Shared Resources

All central processing units (CPUs) in a CRAY C90 series mainframe share the following resources (refer to Figure 2-1):

- Central memory
- I/O section
- Interprocessor communication section
- Real-time clock

Central Memory

Central memory consists of random-access memory (RAM) that is shared by all the CPUs and the I/O section. Each memory word consists of 80 bits: 64 data bits and 16 error-correction bits (check bits). Storage for data and check bits is provided by bipolar complementary metal oxide semiconductor (BiCMOS) chips. In order to improve memory access speed, central memory is divided into multiple banks that can be active simultaneously.

In each CPU, the operating registers, instruction buffers, and exchange package have access to central memory through memory ports. Each CPU has four ports. Each of these ports is 2 words wide, allowing up to eight simultaneous memory references from each CPU. The I/O section shares one port in each CPU.

A CRAY C90 series mainframe central memory uses a single-byte error correction/double-byte error detection (SBCDBD) memory error-correction scheme instead of the single-error correction/double-error detection (SECDED) method used in previous Cray Research machines. SBCDBD ensures that data written into central memory is read with consistent precision. If a single byte (4 bits) of data

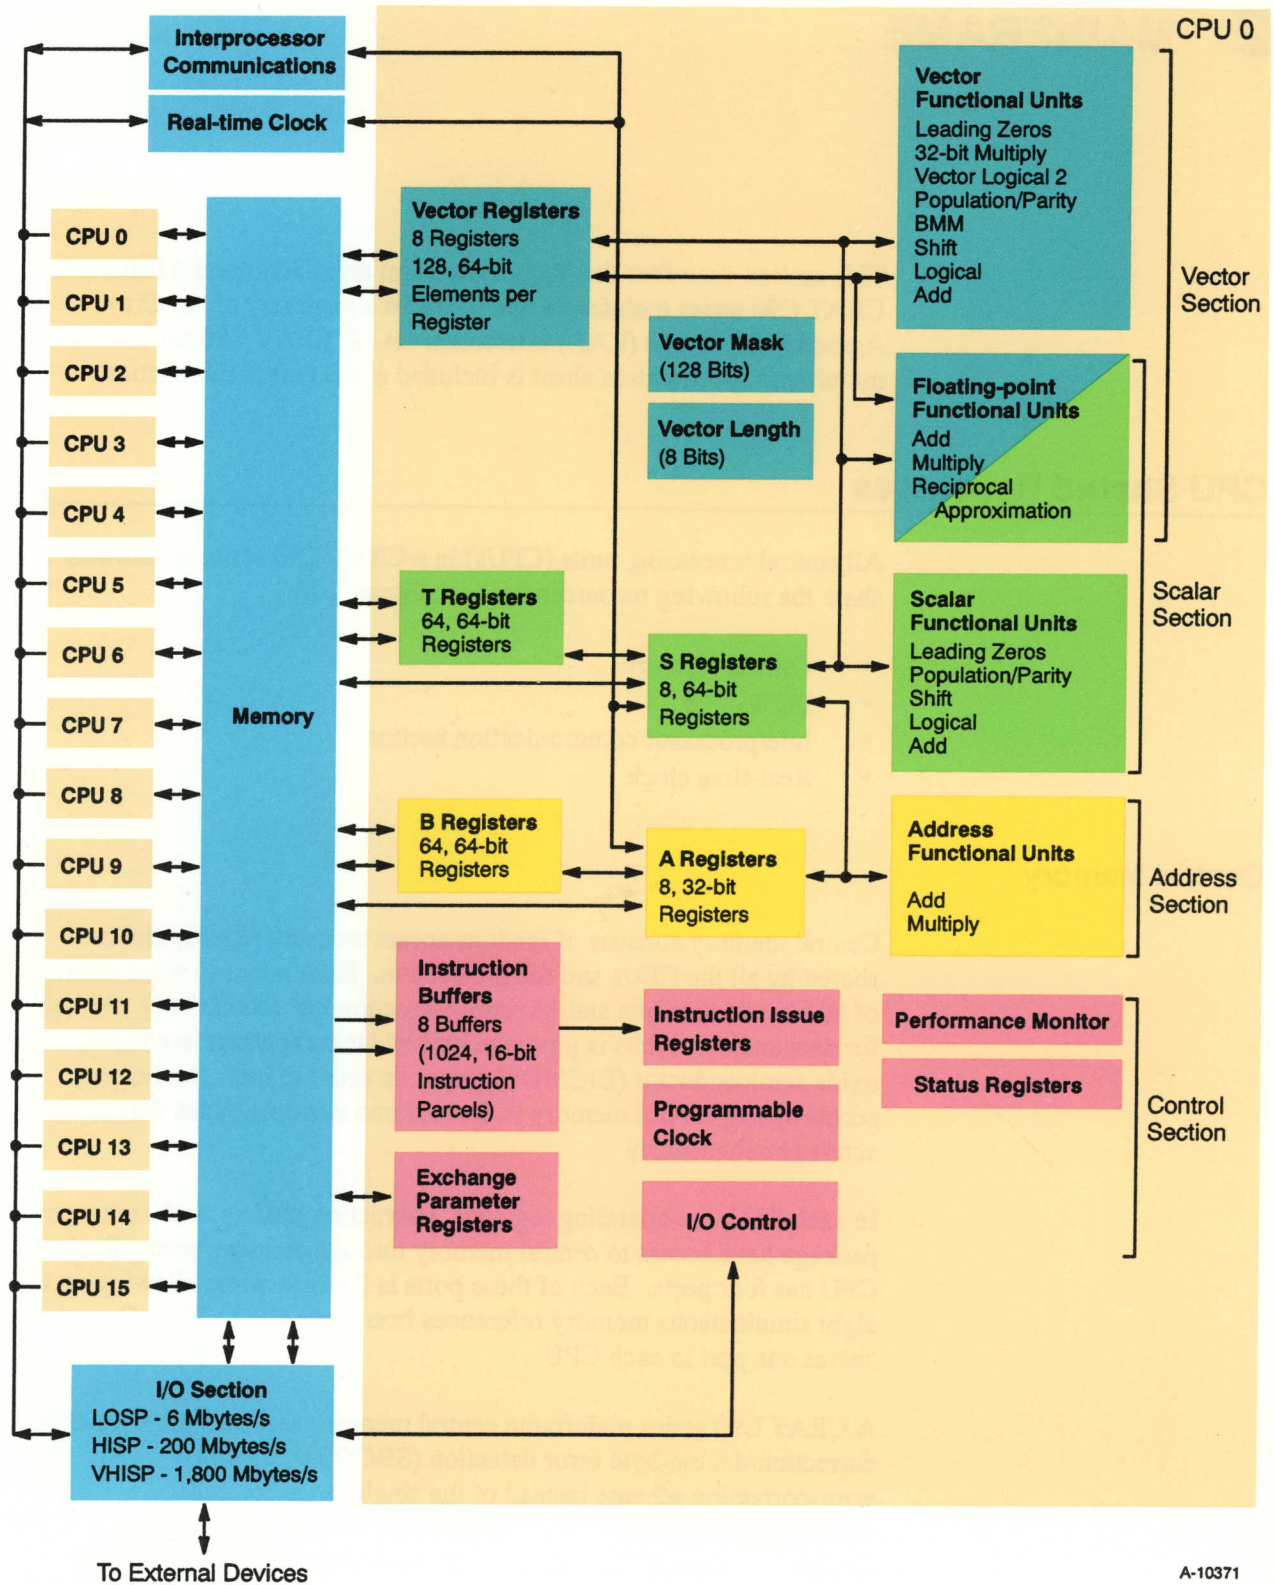


Figure 2-1. CRAY C90 Series Mainframe Block Diagram

is corrupted, the byte is automatically corrected when the word is read from memory. If 2 or more bytes are corrupted, then a double-byte error has occurred and can be detected, but not corrected.

I/O Section

All CPUs share the I/O section of the computer system. The computer system supports three channel types, which are identified by their maximum transfer rates:

- Low-speed (LOSP) channels – 6 Mbytes/s
- High-speed (HISP) channels – 200 Mbytes/s
- Very high-speed (VHISP) channels – 1,800 Mbytes/s

Interprocessor Communication Section

The interprocessor communication section of the computer system contains shared and semaphore registers to pass data and control information between CPUs. It also contains logic to enable any CPU in monitor mode to interrupt any other CPU and cause it to switch from user mode to monitor mode. These features are especially useful in multitasking environments.

The shared and semaphore registers are divided into identical groups called clusters. Each cluster contains eight 32-bit shared address (SB) registers, eight 64-bit shared scalar (ST) registers, and thirty-two 1-bit semaphore (SM) registers. Each CPU is assigned to one cluster, giving it access to the registers in that cluster.

The shared registers provide intermediate storage between CPUs and a way to transfer data between operating registers in different CPUs. One CPU loads a shared register from its address or scalar registers; other CPUs assigned to the same cluster can then transfer the data from the shared register to their own address or scalar registers. Within a CPU, data is transmitted between the SB and address registers and between the ST and scalar registers.

Semaphore (SM) registers allow a CPU to temporarily suspend program operation in order to synchronize operation with other CPUs. Each CPU can set or clear each SM register in its assigned cluster and can perform a test and set instruction on those SM registers. A test and set instruction can result in a CPU holding further execution of instructions until the appropriate SM register is cleared by another CPU assigned to the cluster. Each CPU in the cluster can also transmit all 32 SM registers to or from a scalar register.

Real-time Clock

A CRAY C90 series mainframe has a real-time clock (RTC) that increments synchronously with program execution and may be used to compute the running time for a program in clock periods (CPs). The RTC is a 64-bit counter that increments each CP.

CPU Computation Section

Each CPU is an identical, independent computation section consisting of operating registers, functional units, and an instruction control network (refer again to Figure 2-1). The operating registers and functional units store and process three types of data: address, scalar, and vector.

Address data controls internal operations and consists of information such as memory addresses, register designators and indexes. Address data is stored in the address (A) registers and intermediate address (B) registers and is processed in two dedicated functional units.

Scalar data is any discrete numerical quantity that can be processed in functional units either singly or in operand pairs to produce a single scalar result. Scalar data is stored in the scalar (S) registers and the intermediate scalar (T) registers and is processed in four dedicated functional units. Scalar floating-point data is processed in one of three floating-point functional units; these functional units are also used to process vector floating-point data.

Vector data refers to a set (or vector) of discrete numerical quantities that can be referenced by a single name. Vector data can be processed either singly or in operand pairs in special functional units to produce a vector result. Practically speaking, this means that a single instruction can result in the same operation being performed sequentially on a whole set of operands to produce a set of results. Vector data is stored in the vector (V) registers and is processed in five dedicated functional units. Vector floating-point data is processed in one of three floating-point functional units; these functional units are also used to process scalar floating-point data.

The 32-bit integer product is a vector instruction designed for index calculation. A full-indexing capability is possible throughout central memory in either scalar or vector modes. The index can be positive or negative in either mode. Indexing allows matrix operations in vector mode to be performed on rows or on the diagonal as well as allowing conventional column-oriented operations.

Data flow in a computation section is from central memory to registers and from registers to functional units. Results flow from functional units to registers and from registers to central memory or back to functional

units. Depending on the instruction sequence, data flows along either the scalar or vector path with two exceptions. In some cases, the scalar registers may provide one of the required operands for some vector operations performed in the vector functional units. Also, some scalar functional units return their results to an address register.

The computation section performs integer or floating-point arithmetic operations. Integer arithmetic is performed in two's complement mode; floating-point quantities have signed magnitude representation.

Integer (or fixed-point) operations are integer addition, integer subtraction, and integer multiplication. No integer division instruction is provided; the operation is accomplished through a software algorithm using floating-point hardware.

Floating-point instructions allow addition, subtraction, multiplication, and reciprocal approximation operations. The reciprocal approximation instructions used in conjunction with other instructions enable floating-point division operations.

An optional bit matrix multiply (BMM) functional unit is available. It performs matrix arithmetic operations using the bit matrix multiply algorithm described later in this section. A second vector population/parity/leading zero count functional unit is also added to the CPU when a BMM functional unit is added.

The instruction set includes logical operations for AND, inclusive OR, exclusive OR, exclusive NOR, and mask-controlled merge operations. Shift operations allow the manipulation of either 64-bit or 128-bit operands to produce 64-bit results. With the exception of 32-bit integer arithmetic performed in the A register functional units, most operations are used in vector or scalar instructions.

The following subsections describe the operating registers and their associated functional units.

Operating Registers

Each CPU has three primary and two intermediate sets of operating registers. The primary sets of operating registers are the address (A), scalar (S), and vector (V) registers. These registers are considered primary because functional units and central memory can access them directly.

For the A and S registers, an intermediate level of registers exists. The A registers are supported by the intermediate address (B) registers, and the S registers are supported by the intermediate scalar (T) registers. The B and T registers cannot access the functional units, and serve mainly as a memory buffer for the primary registers. To reduce the number of

memory reference instructions for scalar and address operations, block transfers are possible between the B and T registers and central memory. The V registers do not have associated intermediate registers.

Address Registers

Each CPU contains eight 32-bit A registers. The A registers serve a variety of applications, but are primarily used as address registers for memory references and as index registers. They provide values for shift counts, loop control, and channel I/O operations and receive values of population count and leading zeros count. In address applications, A registers index the base address for scalar memory references and provide both a base address and an address increment for vector memory references.

Each CPU contains 64 B registers; each register is 32 bits wide. The B registers are used as intermediate storage for the A registers. Data is transferred between B registers and central memory, and between A and B registers. Typically, B registers contain data to be referenced repeatedly over a long time, making it unnecessary to retain the data in either A registers or central memory. Examples of data stored in B registers are loop counts, variable array base addresses, and dimensions.

The data stored in B registers are protected with parity bits. When a word is written into a B register, a set of parity bits is generated and stored with the data bits. This set of parity bits is compared to another set that is generated when a word is read out of the B register. An error is indicated when the two sets do not match. Parity errors set the register parity error (RPE) flag in the exchange package if interrupt on register parity error (IRP) mode is set and enabled. They also report the location of the error to the status register.

Scalar Registers

Each CPU contains eight S registers; each register is 64 bits wide. The S registers are the principal scalar registers for a CPU. Scalar registers serve as the source and destination of scalar arithmetic and logical instructions. Scalar registers can also provide an operand for some vector operations.

Each CPU contains 64 T registers; each register is 64 bits wide. The T registers are used as intermediate storage for the S registers. Data is transferred between T registers and central memory, and between T and S registers.

The data stored in T registers are protected with parity bits. When a word is written into a T register, a set of parity bits is generated and stored with the data bits. This set of parity bits is compared to another

set that is generated when a word is read out of the T register. An error is indicated when the two sets do not match. Parity errors set the register parity error (RPE) flag in the exchange package if interrupt on register parity error (IRP) mode is set and enabled. They also report the location of the error to the status register.

Vector Registers

Each CPU contains eight V registers. Each V register contains 128_{10} elements; each element can store 64 bits of data. In vector operations, the 128 elements are processed in two groups, called pipes. One pipe processes the even-numbered elements while the other pipe simultaneously processes the odd-numbered elements. Each pipe is supported by an identical set of functional units.

The effective length of a V register for any operation is controlled by the program-selectable vector length (VL) register. The VL register is an 8-bit register that specifies the number of vector elements processed by the vector instructions. The contents range from 1_8 through 200_8 .

The vector mask (VM) register allows for the logical selection of particular elements of a vector. The VM register is a 128-bit register; each bit corresponds to an element of a vector register. Bit 2^{127} corresponds to element 0 and bit 2^0 corresponds to element 127. The mask is used with vector merge and test instructions to allow operations to be performed on individual vector elements.

V register data is protected with parity bits. When a word is written into a V register, a set of parity bits is generated and stored with the data bits. This set of parity bits is compared to another set that is generated when the word is read out of the V register. An error is indicated when the two sets do not match. Parity errors set the register parity error (RPE) flag in the exchange package if interrupt on register parity error (IRP) mode is set and enabled. They also report the location of the error to the status register.

For more information on vector processing, refer to “Vector Processing” in this section.

Functional Units

Instructions other than simple transfers of data or control operations are performed by specialized hardware known as functional units. Each unit implements an algorithm or a portion of the instruction set. Most functional units have independent logic, and all can operate simultaneously.

All functional units perform their specific operation in a fixed amount of time; delays are impossible once the operands are delivered to the unit. Functional units are fully segmented. This means a new set of operands for unrelated computation can enter a functional unit each CP, even though the functional unit time can be more than 1 CP. Refer to “Pipelining and Segmentation” and “Functional Unit Independence” in this section for more information on pipelining, segmentation, and functional unit independence.

There are four groups of functional units: address, scalar, vector, and floating-point. The address, scalar, and vector functional units operate with one of the primary register types (A, S, and V) to support address, scalar, and vector processing. The floating-point functional units support either scalar or vector operations and accept operands from or deliver results to S or V registers. For timing purposes, central memory can also act as a functional unit for vector operations.

The following subsections define the functions and the instructions executed by each functional unit. Refer to the following sections and subsections for additional information on functional units.

Address Functional Units

Address functional units perform integer arithmetic on operands obtained from A registers and deliver the results to an A register. Integer arithmetic is explained later in this section. The two address functional units are described below.

Address Add Functional Unit

The address add functional unit performs integer addition and subtraction; subtraction is performed by using two’s complement arithmetic. Overflow is not detected.

Address Multiply Functional Unit

The address multiply functional unit forms an integer product from two operands. No rounding is performed, and overflow is not detected. The unit returns only the least significant 32 bits of the product.

Scalar Functional Units

Scalar functional units perform operations on operands obtained from S registers and usually deliver the results to an S register. The exception is the population/parity/leading zero count functional unit, which delivers its result to an A register.

Four functional units are exclusively associated with scalar operations and are described below. Three floating-point functional units are used for both scalar and vector operations. Refer to “Floating-point Functional Units” in this section for more information on these units.

Scalar Add Functional Unit

The scalar add functional unit performs integer addition and subtraction; subtraction is performed by using two’s complement arithmetic. Overflow is not detected.

Scalar Logical Functional Unit

The scalar logical functional unit performs bit-by-bit manipulation of quantities obtained from S registers.

Scalar Shift Functional Unit

The scalar shift functional unit shifts the entire contents of an S register (single shift) or shifts the contents of two concatenated S registers (double shift) into a single resultant S register. Single shifts are end-off with zero fill, while double shifts can be circular fill. Shift counts are obtained from an A register or from a field of the instruction.

Scalar Population/Parity/Leading Zero Functional Unit

The scalar population/parity/leading zero count functional unit counts the number of 1 bits in an operand obtained from an S register and then, depending on the instruction issued, returns the count either as a population or population parity count to an A register. For the leading zero function, the unit counts the number of 0 bits preceding the first 1 bit in an operand obtained from an S register and returns the count to an A register.

Vector Functional Units

There are two parallel sets of vector functional units referred to as pipe 0 and pipe 1. Pipe 0 processes the even-numbered elements of a vector, while pipe 1 processes the odd-numbered elements. This duplication of functional units allows two pairs of elements to be processed at the same time and increases the efficiency of the vector processing operations.

Most vector functional units perform operations on operands obtained from one or two vector registers or from a vector register and an S register. The shift, population/parity, and leading zero functional units require only one operand. Results from a vector functional unit are delivered to a V register.

The functional units described in this section are used exclusively for vector operations. Three functional units are associated with both vector operations and scalar operations. Refer to “Floating-point Functional Units” in this section for more information on these functional units.

Vector Add Functional Unit

The vector add functional unit performs integer addition and subtraction for a vector operation and delivers the results to elements of a V register. The subtraction operation uses two’s complement arithmetic. Overflow is not detected.

Vector Shift Functional Unit

The vector shift functional unit shifts the entire contents of a vector register element (single shift) or the value formed from two consecutive elements of a V register (double shift). Shift counts are obtained from an A register and are end-off with zero fill.

Full Vector Logical Functional Unit

The full vector logical functional unit performs a bit-by-bit manipulation of specified quantities for specific instructions. The full vector logical functional unit also performs vector register merge, compressed index, and the logical operations associated with the vector mask instructions.

Second Vector Logical Functional Unit

The second vector logical functional unit, when enabled, performs the same type of bit-by-bit manipulations as the full vector logical functional unit, but not for all instructions. The second vector logical functional unit cannot perform vector register merge, compressed index, and the logical operations associated with the vector mask instructions. A bit in the exchange package enables or disables the second vector logical functional unit.

Vector Population/Parity/Leading Zero Functional Unit

The vector population/parity/leading zero count functional unit performs population counts, parity checks, and leading zero counts for vector operations. These operations are identical to those performed in the scalar population/parity/leading zero count functional unit, except that the operands are the elements of a V register, and the results are returned to a V register.

Second Vector Population/Parity/Leading Zero Functional Unit

The optional second vector population/parity/leading zero functional unit is included with CPUs that have the optional bit matrix multiply (BMM) functional unit. The second vector population/parity/leading zero functional unit enables the CPU to chain BMM and population/parity/leading zero operations. If the first population/parity/leading zero functional unit is busy at instruction issue time, the operation is sent to the second population/parity/leading zero functional unit.

Bit Matrix Multiply Functional Unit

The optional BMM functional unit performs a logical multiplication of two square matrices, resulting in a single bit for each pair of elements of the matrices. The matrices, which are held in the vector registers, vary in size from 1×1 to 64×64 . The size of the matrix is specified by the contents of the vector length (VL) register.

In addition to performing full 64×64 matrix multiply operations on the contents of two vector registers, the BMM functional unit can perform a scalar-vector multiply on the contents of a vector register and a scalar register and store the result in an S register.

Floating-point Functional Units

There are two parallel sets of floating-point functional units, with each set containing three functional units. These floating-point functional units perform floating-point arithmetic for both scalar and vector operations. The vector registers use both sets of functional units; one set processes the even-numbered elements, while the other set processes the odd-numbered elements. For an operation involving only scalar operands, only one set of floating-point functional units is used.

When executing most vector instructions, operands are obtained from pairs of V registers, or from an S register and a V register, and results are delivered to a vector register. When a floating-point functional unit

used for a vector operation, the general description of vector functional units applies. When executing a scalar instruction, operands are obtained solely from S registers, and results are delivered to an S register.

Floating-point Add Functional Unit

The floating-point add functional unit performs addition and subtraction of operands in floating-point format. The result is normalized even when operands are unnormalized. The floating-point add functional unit detects overflow and underflow conditions; only overflow conditions are flagged.

Floating-point Multiply Functional Unit

The floating-point multiply functional unit performs full- and half-precision multiplication of operands in floating-point format. The half-precision product is rounded; the full-precision product can be rounded or not rounded. This functional unit also generates a 32-bit integer product.

Input operands must be normalized; the floating-point multiply functional unit delivers a normalized result only if both input operands are normalized. The floating-point multiply functional unit detects overflow and underflow conditions; only overflow conditions are flagged.

The floating-point multiply functional unit recognizes both operands with zero exponents as a special case and performs an integer multiply operation. The result is considered an integer product, is not normalized, and is not considered out of range.

Reciprocal Approximation Functional Unit

The reciprocal approximation functional unit finds the approximate reciprocal of an operand in floating-point format. The input operand must be normalized; the floating-point reciprocal approximation functional unit delivers a correct result only if the input operand is normalized. The high-order bit of the coefficient is not tested, but is assumed to be a 1. The floating-point reciprocal approximation functional unit detects overflow and underflow conditions; both conditions are flagged.

Functional Unit Operations

Functional units in a CPU perform logical operations, integer arithmetic, and floating-point arithmetic. Integer arithmetic and floating-point arithmetic are performed in two's complement. The following subsections explain the logical operations, the integer arithmetic, and the floating-point arithmetic used by a CRAY C90 series mainframe.

Logical Operations

Scalar and vector logical functional units perform bit-by-bit manipulation of 64-bit quantities. Instructions are provided for forming logical products, sums, exclusive ORs, equivalences, and merges.

A logical product is the AND function, which is shown in the following example:

$$\begin{array}{r} \text{Operand 1: } 1\ 0\ 1\ 0 \\ \text{Operand 2: } \underline{1\ 1\ 0\ 0} \\ \text{Result: } \quad \underline{1\ 0\ 0\ 0} \end{array}$$

A logical sum is the inclusive OR function, which is shown in the following example:

$$\begin{array}{r} \text{Operand 1: } 1\ 0\ 1\ 0 \\ \text{Operand 2: } \underline{1\ 1\ 0\ 0} \\ \text{Result : } \quad \underline{1\ 1\ 1\ 0} \end{array}$$

A logical exclusive OR function is shown in the following example:

$$\begin{array}{r} \text{Operand 1: } 1\ 0\ 1\ 0 \\ \text{Operand 2: } \underline{1\ 1\ 0\ 0} \\ \text{Result: } \quad \underline{0\ 1\ 1\ 0} \end{array}$$

A logical equivalence is the exclusive NOR function, which is shown in the following example:

$$\begin{array}{r} \text{Operand 1: } 1\ 0\ 1\ 0 \\ \text{Operand 2: } \underline{1\ 1\ 0\ 0} \\ \text{Result: } \quad \underline{1\ 0\ 0\ 1} \end{array}$$

The merge operation uses two operands and a mask to produce results. The bits of operand 1 are transmitted to the result when the mask bit is a 1. The bits of operand 2 are transmitted to the result when the mask bit is a 0. The following example shows a merge operation:

```

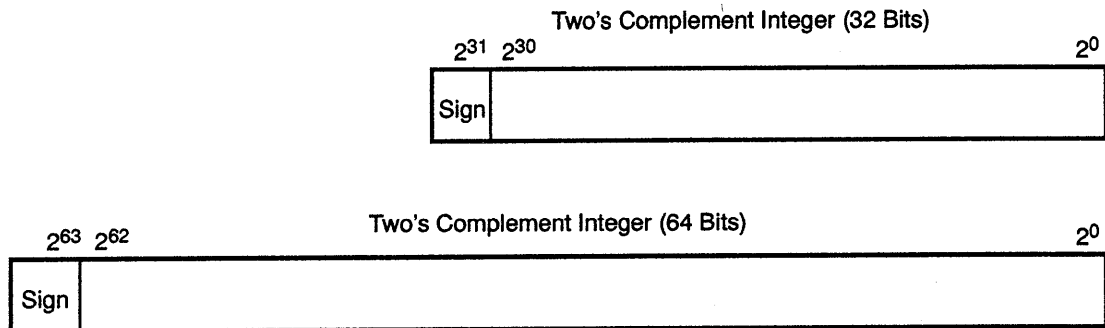
Operand 1: 1 0 1 0 1 0 1 0
Operand 2: 1 1 0 0 1 1 0 0
Mask:      1 1 1 1 0 0 0 0
Result:    1 0 1 0 1 1 0 0
  
```

Integer Arithmetic

All integers, whether 32 or 64 bits long, are represented in the registers as shown in Figure 2-2. The address add and address multiply functional units perform 32-bit arithmetic. The scalar add and vector add functional units perform 64-bit arithmetic.

Two scalar (64-bit) integer operands are multiplied using the floating-point multiply instruction and one of two multiplication methods. The method used depends on the magnitude of the operands and the number of bits available to contain the product. The following paragraphs explain the 24-bit integer multiply operation and the method used for operands greater than 24 bits.

The floating-point multiply functional unit recognizes a condition in which both operands have zero exponents as a special case. This case is treated as an integer multiplication operation, and a complete multiplication operation is performed with no truncation as long as the total number of bits in the two operands does not exceed 48 bit positions. To multiply two integer numbers together, set each operand's exponent (bits 2^{62} through 2^{48}) equal to 0 and place each 24-bit integer value in bit positions 2^{47} through 2^{24} of the operand's coefficient field. To ensure accuracy, the least significant 24 bits must be 0's.




A-10372

Figure 2-2. Integer Data Formats

When the floating-point multiply functional unit performs the operation, it returns the 48 high-order bits of the product as the result coefficient and leaves the exponent field as 0. The result is a 48-bit quantity in bit positions 2^{47} through 2^0 ; no normalization shift of the result is performed. If the 24 least significant bits of the operand coefficients were nonzero, the 48 low-order bits of the product could be nonzero and could generate a carry into the least significant of the 48 high-order bits returned, causing the result to be one larger than expected.

As shown in Figure 2-3, if operand 1 is 4 and operand 2 is 6, a 48-bit result of 30g is produced. Bit 2^{63} follows the rules for multiplying signs, and the result is a signed-magnitude integer. An exclusive OR function on bits 2^{63} of operands 1 and 2 is performed to derive the sign of the result.



A-10373

Figure 2-3. 24-bit Integer Multiply Performed in a Floating-point Multiply Functional Unit

The format of integers expected by both the hardware and software is two's complement, not signed-magnitude; therefore, negative products must be converted to two's complement form.

The second multiplication method is used when the operands are more than 24 bits long; multiplication is done by software, which forms multiple partial products and then shifts and adds the partial products.

A second integer multiplication operation performs a 32-bit multiplication operation on the S_j operand and the V_k operand and puts the result in the V_i register. The operands must be shifted left before the operation begins. The S_j operand must be shifted left 31_{10} places, leaving the operand in bit positions 2^{62} through 2^{31} ; bit positions 2^{30} through 2^0 must be equal to 0 to ensure accuracy (refer to Figure 2-4). The V_k operand must be shifted left 16_{10} places, leaving the operand in bit positions 2^{47} through 2^{16} ; bit positions 2^{15} through 2^0 must be equal

to 0 to ensure accuracy. Bits 2^{63} through 2^{48} are zero filled. The result of the multiply is right justified into positions 2^{31} through 2^0 , and positions 2^{32} through 2^{63} are zero filled.




Figure 2-4. 32-bit Integer Multiply Performed in a Floating-point Multiply Functional Unit

Although no integer division operation is provided, integer division can be carried out by converting the numbers to the floating-point format and then using the floating-point functional units. For more information on integer division, refer to “Floating-point Division Algorithm” in this section.

Floating-point Arithmetic

The scalar and vector instructions use floating-point arithmetic. The following subsections explain floating-point arithmetic.

Floating-point Data Format

Floating-point numbers are represented in a standard format throughout the CPU; this format is shown in Figure 2-5. The format has three fields: coefficient sign, exponent, and coefficient.




Figure 2-5. Floating-point Data Format

This format is a packed representation of a binary coefficient and an exponent (power of two). The coefficient sign is located in bit position 2^{63} and is separated from the rest of the coefficient. If this bit is equal to 0, the coefficient is positive; if this bit is equal to 1, the coefficient is negative.

The exponent is represented as a biased integer number in bit positions 2^{62} through 2^{48} ; each exponent is biased by 40000_8 . Figure 2-6 shows the biased and unbiased exponent ranges. Bit 2^{61} is the sign of the exponent; a 0 indicates a positive exponent, and a 1 indicates a negative exponent. Bit 2^{62} is the bias of the exponent. The floating-point format of the system allows the accurate expression of numbers to about 15 decimal digits in the approximate range of 10^{-2466} through 10^{+2466} .

The coefficient is a 48-bit signed fraction; the sign of the coefficient is located in bit position 2^{63} . Because the coefficient is in signed-magnitude format, it is not complemented for negative values.




Figure 2-6. Biased and Unbiased Exponent Ranges

Figure 2-7 and the following steps show the relation between the biased exponent and the coefficient. The following steps show how to convert a floating-point number to its decimal equivalent.




Figure 2-7. Internal Representation of a Floating-point Number

1. Subtract the bias from the exponent to get the integer value of the exponent:

$$\begin{array}{r} 40011_8 \\ -40000_8 \\ \hline 11_8 = 9_{10} \end{array}$$

2. Multiply the normalized coefficient by the power of 2 indicated in the exponent to get the result:

$$0.5634_8 \times 2^9 = 563.40_8 = 371.5_{10}$$

A zero value or an underflow result is not biased and is represented as a word of all 0's. A negative 0 is not generated by any floating-point functional unit, except in the case in which a negative 0 is one operand going into the floating-point multiply or floating-point add functional unit.

Normalized Floating-point Numbers

A nonzero floating-point number is normalized if the most significant bit of the coefficient (bit 2^{47}) is nonzero. This condition implies that the coefficient has been shifted as far left as possible and that the exponent has been adjusted accordingly; therefore, a normalized floating-point number has no leading 0's in its coefficient. The exception is a normalized floating-point 0, which is all 0's.

Anytime an integer is converted to a floating-point number, normalize the result before using it in a floating-point operation. Normalization is accomplished by adding the unnormalized floating-point operand to 0.

The reciprocal approximation functional unit must use normalized numbers to produce correct results. Using unnormalized numbers produces inaccurate results.

The floating-point multiply functional unit does not require the use of normalized numbers to get correct results. However, more accurate results occur when normalized numbers are used.

The floating-point add functional unit does not require normalized numbers to get correct results. The floating-point add functional unit does, however, automatically normalize all its results; unnormalized floating-point numbers may be routed through this functional unit to take advantage of this process.

Floating-point Range Errors

To ensure that the limits of the functional units are not exceeded, a range check is performed for overflow and underflow conditions on the exponent of each floating-point number coming into the functional unit. Bits 2^{61} and 2^{62} are checked; if both bits are equal to 1, the exponent is equal to or greater than 60000_8 , and an overflow condition is detected.

When an overflow condition is detected, an interrupt occurs only if the interrupt-on-floating-point error (IFP) mode is set and enabled. In this case, the floating-point error (FPE) flag is set, causing an exchange sequence to occur. The IFP mode can be set or cleared by a user mode program.

When an overflow condition occurs, the value returned to the result register depends on the functional unit used. For the floating-point add and floating-point multiply functional units, the calculated coefficient, together with a forced exponent of 60000_8 , is sent to the result register. For the reciprocal approximation functional unit, the returned result is the same except that bit 2^{47} of the coefficient is set to 0. Refer to Figure 2-8 and Figure 2-9.





Figure 2-8. Floating-point Add and Floating-point Multiply Range Errors

To check for an underflow condition in the floating-point functional units, bits 2^{61} and 2^{62} are checked; if both are equal to 0, then the exponent is less than or equal to 17777_8 , and an underflow condition is detected.

If an underflow condition is detected in the floating-point add or floating-point multiply functional unit, no fault is generated, and the word returned from the functional unit is all 0 bits. Refer to Figure 2-8. The floating-point multiply functional unit will not detect an underflow condition if both exponents equal 0; instead an integer multiply operation is performed.

Because the underflow condition of the result generated by the floating-point add functional unit is tested before the result is normalized, the normalized result can have a valid exponent as low as 17721_8 . This occurs when the unnormalized result has an exponent of 20000_8 and a coefficient of 1. In this case, no underflow is detected, and the calculated result is sent to the result register.

An underflow condition is detected in the reciprocal approximation functional unit if either of the incoming operands has an exponent less than or equal to 20001_8 . If this condition occurs, the FPE flag will set only if IFP mode is set and enabled. The calculated coefficient, with bit 2^{47} set to 0, together with a forced exponent of 60000_8 , is sent to the result register. Refer to Figure 2-9.



A-10379

Figure 2-9. Floating-point Reciprocal Approximation Range Errors

Floating-point Addition Algorithm

Floating-point addition or subtraction is performed in a 49-bit register to allow for a sum that might carry into an additional bit position. The algorithm performs three operations: equalizing exponents, adding coefficients, and normalizing results.

To equalize the exponents, only the larger of the two exponents is retained. The coefficient of the smaller exponent is shifted right by the difference of the two exponents. Bits shifted out of the register are lost; no roundup occurs. Because the coefficient is only 48 bits long, any shift beyond 48 bits causes the smaller coefficient to become 0.

After the two coefficients are equalized, they are added together. Two conditions are analyzed to determine whether an addition or subtraction operation occurs. The two conditions are the sign bits of the two coefficients and the type of instruction (an add or subtract) issued. The following list shows how the operation is determined.

- If the sign bits are equal and an add instruction is issued, an addition operation is performed.
- If the sign bits are not equal and an add instruction is issued, a subtraction operation is performed.
- If the sign bits are equal and a subtract instruction is issued, a subtraction operation is performed.
- If the sign bits are not equal and a subtract instruction is issued, an addition operation is performed.

The last operation performed normalizes the results. To normalize the result, the coefficient is shifted left by the number of leading 0's (the coefficient is normalized when bit 2^{47} is a 1). The exponent must also be decremented accordingly. If a carry across the binary point occurs during an addition operation, the coefficient is shifted right by 1 and the exponent increases by 1.

The normalization feature of the floating-point add functional unit is used to normalize any floating-point number. The number is simply paired with a zero operand and sent through the floating-point add functional unit.

A range check is performed on the result of all additions; refer to "Floating-point Range Errors" earlier in this subsection for more information on how the result is checked.

Floating-point Multiplication Algorithm

The floating-point multiply functional unit receives two 48-bit floating-point coefficients from either an S or V register as input. Multiplication is commutative, that is, $A \times B = B \times A$. The signs of the two operands are combined by an exclusive OR function, the exponents are added together, and the two 48-bit coefficients are multiplied together. Multiplying the 48-bit coefficients produces a product of either 95 or 96 bits. A 96-bit product is normalized as it is generated, but a 95-bit product requires a left shift of 1 to generate the final normalized coefficient. If a shift occurs, the final exponent is reduced by 1 to reflect the shift.

Because the result register (an S or V register) can hold only 48 bits in the coefficient, only the upper 48 bits of the 96-bit result are used. Some of the lower 48 bits are never generated. To adjust for this truncation, a constant is unconditionally added to the product. The average value of this truncation is 9.25×2^{-56} , which is determined by adding all carries produced by all possible combinations that could be truncated and dividing the sum by the number of possible combinations. Nine carries are inserted at bit position 2^{-56} to compensate for the truncated bits.

If the truncated bits are not compensated for, the resulting coefficient is 1 bit position smaller than expected. With compensation, the resulting coefficient ranges from 1 too large to 1 too small in the 2^{-48} bit position, with approximately 99% of the values having zero deviation from what would have been generated had a full 96-bit product been present. Rounding is optional, but truncation compensation is not. The rounding method used adds a constant so that it is 50% high (0.25×2^{-48} ; high) 38% of the time, and 25% low (0.125×2^{-48} ; low) 62% of the time, resulting in a near-zero average rounding error. In a full-precision rounded multiplication operation, 2 round bits are entered into the summation at bit positions 2^{-50} and 2^{-51} and allowed to propagate.

For a half-precision multiplication operation, round bits are entered into the summation at bit positions 2^{-32} and 2^{-31} . A carry resulting from this entry is allowed to propagate upward, and the 29 most significant bits of the normalized result are transmitted back.

The result variations caused by this truncation and rounding are in the following ranges:

$$\begin{aligned} & -0.23 \times 2^{-48} \text{ to } +0.57 \times 2^{-48} \\ & \text{or} \\ & -8.17 \times 10^{-16} \text{ to } +20.25 \times 10^{-16} \end{aligned}$$

With a 96-bit product and rounding equal to one-half the least significant bit, the following result variation is expected:

$$-0.5 \times 2^{-48} \text{ to } +0.5 \times 2^{-48}$$

Floating-point Division Algorithm

A CRAY C90 series mainframe does not have a single functional unit dedicated to the division operation. Rather, the floating-point multiply and reciprocal approximation functional units together carry out the algorithm. The following paragraphs explain the algorithm and how it is used in the functional units.

Finding the quotient of two floating-point numbers involves two steps. For example, to find the quotient A/B , first the B operand is sent through the reciprocal approximation functional unit to obtain its reciprocal, $1/B$. Then, this result along with the A operand is sent to the floating-point multiply functional unit to obtain the product $A \times 1/B$.

The reciprocal approximation functional unit uses an application of Newton's method for approximating the real root of an arbitrary equation, $F(x) = 0$, to find reciprocals. Refer to Figure 2-10.

To find the reciprocal, the equation $F(x) = 1/x - B = 0$ must be solved. To do this, a number, A , must be found so that $F(A) = 1/A - B = 0$. That is, the number A is the root of the equation $1/x - B = 0$. The method requires an initial approximation (or guess, which is shown as x_0 in Figure 2-10) sufficiently close to the true root (which is shown as x_t in Figure 2-10). The initial approximation, x_0 , is then used to obtain a better approximation; this is done by drawing a tangent line (line 1 in Figure 2-10) to the graph of $y = F(x)$ at the point $[x_0, F(x_0)]$. The x -intercept of this tangent line becomes the second approximation, x_1 . This process is repeated, using tangent line 2 to obtain x_2 , and so on.

The following iteration equation is derived from this process:

$$x_{(i+1)} = 2x_i - x_i^2 B = x_i (2 - x_i B)$$

In the equation, $x_{(i+1)}$ is the next iteration, x_i is the current iteration, and B is the divisor. Each $x_{(i+1)}$ is a better approximation than x_i to the true value, x_t . The exact answer is generally not obtained at once because the correction term is not exact. The operation is repeated until the answer becomes sufficiently close for practical use.

The reciprocal approximation functional unit calculates the first two iterations, while the floating-point multiply functional unit calculates the third iteration. The third iteration uses a special instruction within the floating-point multiply functional unit to calculate the correction term. This iteration is used to increase accuracy of the reciprocal approximation functional unit's answer to full precision. The floating-point multiply functional unit can provide both full- and half-precision results.

The reciprocal iteration is designed for use once with each half-precision reciprocal generated. If the third iteration (the iteration performed by the floating-point multiply functional unit) results in an exact reciprocal, or if an exact reciprocal is generated by some other method, performing another iteration results in an incorrect final reciprocal. A fourth iteration should not be done.

The following example shows how the floating-point multiply functional unit provides a full-precision result, computing the value of $S1/S2$.

| <u>Step</u> | <u>Operation</u> | <u>Unit</u> |
|-------------|------------------------|---|
| 1 | $S3 = 1/S2$ | Reciprocal approximation functional unit |
| 2 | $S4 = [2 - (S3 * S2)]$ | Floating-point multiply functional unit |
| 3 | $S5 = S4 * S3$ | Floating-point multiply functional unit using full-precision; S5 now equals $1/S2$ to 48-bit accuracy |
| 4 | $S6 = S5 * S1$ | Floating-point multiply functional unit using full-precision rounding |

The reciprocal approximation in Step 1 is correct to 30 bits. By Step 3, it is accurate to 48 bits. This iteration answer is applied as an operand in a full-precision rounded multiplication operation (Step 4) to obtain a quotient accurate to 48 bits. Additional iterations may produce erroneous results.

Where 29 bits of accuracy are sufficient, the reciprocal approximation instruction is used with the half-precision multiply to produce a half-precision quotient in only two operations, as shown in the following example.

| <u>Step</u> | <u>Operation</u> | <u>Unit</u> |
|-------------|------------------|---|
| 1 | $S3 = 1/S2$ | Reciprocal approximation functional unit |
| 2 | $S6 = S1 * S3$ | Floating-point multiply functional unit in half-precision |

The 19 low-order bits of the half-precision multiply results are returned as 0's with rounding applied to the low-order bit of the 29-bit result.

The following is another method of performing the division operation:

| <u>Step</u> | <u>Operation</u> | <u>Unit</u> |
|-------------|------------------------|--|
| 1 | $S3 = 1/S2$ | Reciprocal approximation functional unit |
| 2 | $S5 = S1 * S3$ | Floating-point multiply functional unit |
| 3 | $S4 = [2 - (S3 * S2)]$ | Floating-point multiply functional unit |
| 4 | $S6 = S4 * S5$ | Floating-point multiply functional unit |

With this method, the correction to reach a full-precision reciprocal is done after the numerator is multiplied by the half-precision reciprocal rather than before the multiplication.

The coefficient of the reciprocal produced by this alternative method can differ by as much as 2×2^{-48} from the first method described for generating full-precision reciprocals. This difference can occur because one method can round up as much as twice, while the other method may not round at all. The first rounding can occur while the correction is generated, and the second rounding can occur when the final quotient is produced. Therefore, the reciprocals should be compared using the same method each time they are generated. Cray Fortran CFT and CFT77 use a consistent method to ensure that the reciprocals of numbers are always the same.

Double-precision Numbers

The CPU does not provide special hardware for performing double- or multiple-precision operations. Double-precision computations with 95-bit accuracy are available through software routines provided by Cray Research.

Bit Matrix Multiply Arithmetic

The vector matrix multiply functional unit performs a logical multiplication of two square bit matrices of equal size. The size varies from 1 x 1 to 64 x 64. Because the matrices must be square, a vector length of 20 (VL = 20) indicates a 20 x 20 bit matrix. In the case of a matrix of less than 64 bits, the contents of the matrix must be left-justified and zero-filled in the unused bit positions. For example, in a 20 x 20 matrix the contents of elements 0 through 19 must be left-justified and zero-filled in bit positions 2⁰ through 2⁴³. Data stored in elements 20 through 63 is not used, and the functional unit treats this data as a “don’t care” condition. Refer to Figure 2-11.

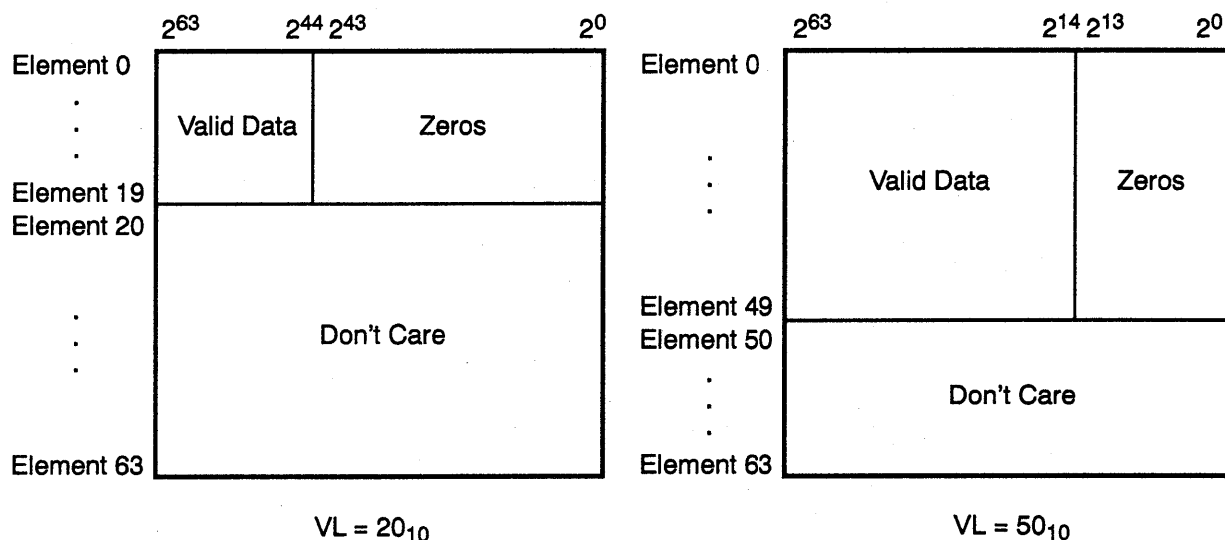


Figure 2-11. Vector Storage of a Bit Matrix

All matrices are stored in vector registers. Each vector element holds the contents of a separate row of the matrix; each bit of the element is a column entry of the respective row that the element represents. Throughout this subsection, the terms *row* and *column* are used when referring to matrices, and the terms *element* and *bit* are used when referring to vector registers.

The matrix multiply operation is defined in the following example, where:

A and B are two $n \times n$ bit matrices, where $n \leq 64$

Matrix B^t is the transpose of matrix B

Matrix C is the product of matrix A and matrix B^t

Refer to Figure 2-12 through Figure 2-14 for examples of these matrix operations.

The entries in each matrix are represented by lowercase letters with two subscripts. The first subscript denotes the row and the second subscript denotes the column in which the entry is located. For example, a_{23} represents the entry in row 2, column 3 of the A matrix.

$$A = \begin{pmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \dots & a_{3n} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \vdots & \dots & \vdots \\ a_{n1} & a_{n2} & a_{n3} & \dots & a_{nn} \end{pmatrix} \quad B = \begin{pmatrix} b_{11} & b_{12} & b_{13} & \dots & b_{1n} \\ b_{21} & b_{22} & b_{23} & \dots & b_{2n} \\ b_{31} & b_{32} & b_{33} & \dots & b_{3n} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \vdots & \dots & \vdots \\ b_{n1} & b_{n2} & b_{n3} & \dots & b_{nn} \end{pmatrix}$$

Figure 2-12. Matrix A and Matrix B

$$B = \begin{pmatrix} b_{11} & b_{12} & b_{13} & \dots & b_{1n} \\ b_{21} & b_{22} & b_{23} & \dots & b_{2n} \\ b_{31} & b_{32} & b_{33} & \dots & b_{3n} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \vdots & \dots & \vdots \\ b_{n1} & b_{n2} & b_{n3} & \dots & b_{nn} \end{pmatrix} \quad B^t = \begin{pmatrix} b_{11} & b_{21} & b_{31} & \dots & b_{n1} \\ b_{12} & b_{22} & b_{32} & \dots & b_{n2} \\ b_{13} & b_{23} & b_{33} & \dots & b_{n3} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \vdots & \dots & \vdots \\ b_{1n} & b_{2n} & b_{3n} & \dots & b_{nn} \end{pmatrix}$$

Figure 2-13. Matrix B and Matrix B^t

$$\begin{matrix}
 A & B^t = & * & = & C \\
 \left| \begin{array}{cccccc}
 a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\
 a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\
 a_{31} & a_{32} & a_{33} & \dots & a_{3n} \\
 \vdots & \vdots & \vdots & & \vdots \\
 \vdots & \vdots & \vdots & & \vdots \\
 a_{n1} & a_{n2} & a_{n3} & \dots & a_{nn}
 \end{array} \right| & & & & \left| \begin{array}{cccccc}
 b_{11} & b_{21} & b_{31} & \dots & b_{n1} \\
 b_{12} & b_{22} & b_{32} & \dots & b_{n2} \\
 b_{13} & b_{23} & b_{33} & \dots & b_{n3} \\
 \vdots & \vdots & \vdots & & \vdots \\
 \vdots & \vdots & \vdots & & \vdots \\
 b_{1n} & b_{2n} & b_{3n} & \dots & b_{nn}
 \end{array} \right| & & & & \left| \begin{array}{cccccc}
 c_{11} & c_{21} & c_{13} & \dots & c_{1n} \\
 c_{21} & c_{22} & c_{23} & \dots & c_{2n} \\
 c_{31} & c_{32} & c_{33} & \dots & c_{3n} \\
 \vdots & \vdots & \vdots & & \vdots \\
 \vdots & \vdots & \vdots & & \vdots \\
 c_{n1} & c_{n2} & c_{n3} & \dots & c_{nn}
 \end{array} \right|
 \end{matrix}$$

Figure 2-14. Matrix C

The entries of the C matrix are determined from the following rules:

$$\begin{aligned}
 c_{11} &= a_{11}b_{11} \oplus a_{12}b_{12} \oplus a_{13}b_{13} \oplus \dots \oplus a_{1n}b_{1n} \\
 c_{12} &= a_{11}b_{21} \oplus a_{12}b_{22} \oplus a_{13}b_{23} \oplus \dots \oplus a_{1n}b_{2n} \\
 &\vdots \\
 &\vdots \\
 c_{32} &= a_{31}b_{21} \oplus a_{32}b_{22} \oplus a_{33}b_{23} \oplus \dots \oplus a_{3n}b_{2n} \\
 &\vdots \\
 &\vdots
 \end{aligned}$$

The \oplus sign indicates an exclusive OR operation. The expression $a_{11}b_{21}$ represents bit AND operations. In other words, to obtain the entry c_{rc} multiply each bit in row r of A by its corresponding bit in column c of B^t and form the exclusive OR of the products.

A sequence of steps is used to perform the multiplication.

First, the program issues a 1740j4 instruction to load the functional unit with the B matrix stored in the V_j register. When the instruction issues, the functional unit first fills the B storage area with zeros for the unused elements, and then reads the rows of the B matrix one per clock period, storing them as the columns of B^t .

Second, the program issues a 174ij6 instruction to multiply the rows of matrix A stored in vector register V_j with the B^t matrix to produce the result matrix C. The rows of matrix A are streamed through the functional unit one per clock period. As each row of matrix A passes through the unit, it is simultaneously multiplied by all columns of B^t , using the exclusive OR operation ($c_{11} = a_{11}b_{11} \oplus a_{12}b_{12} \oplus a_{13}b_{13} \oplus \dots \oplus a_{1n}b_{1n}$) to generate a single row of the result matrix C.

CPU Control Section

Each central processing unit (CPU) is assigned tasks and is controlled in the execution of those tasks through exchange sequences, fetch sequences, and issue sequences. These three sequences are closely related. For an initial deadstart program or a new program to run, an exchange sequence must occur. This sequence of steps sets several important parameters of the program in the CPU and may initialize some of the CPU's operating registers. A fetch sequence begins immediately after the exchange sequence and transfers a block of instructions from memory to an instruction buffer. The issue sequence then selects the instruction indicated by the program address (P) register, decodes it, determines whether the required registers or functional units are available, and if so, allows the instruction to be executed.

As the instruction executes, the P register increments, causing new instructions to be selected from an instruction buffer and to move through the issue sequence. When a desired instruction is not currently in an instruction buffer, another fetch sequence occurs, retrieving another block of instructions from memory. This overall process continues until either the program terminates or is interrupted, at which time another exchange sequence occurs and the entire process begins again.

The following subsections describe the exchange mechanism, the instruction fetch sequence, and the instruction issue sequence unique to each CPU. The programmable clock, the status register, and the performance monitor are also briefly described.

Exchange Mechanism

Each CPU uses an exchange mechanism for switching instruction execution from program to program. This exchange mechanism transfers blocks of program parameters (known as exchange packages) during a CPU operation, which is referred to as an exchange sequence.

The following subsections describe the contents of the exchange package and explain the exchange sequence in more detail.

Exchange Sequence

The exchange sequence moves the contents of an inactive exchange package from memory into the operating registers. Simultaneously, the exchange sequence retrieves data from the operating registers, uses it to construct the active exchange package, and then moves this exchange package back into memory. This swapping operation occurs in a fixed sequence when all computational activity associated with the active exchange package stops.

The exchange sequence involves 16 memory read references and 16 memory write references. A single 16-word block of memory data is used as the source of the inactive exchange package and the destination of the active exchange package. Word 0 of the active exchange package is swapped with word 0 of the inactive exchange package. The location of this block of data is specified by the contents of the XA register and is part of the active exchange package.

Exchange Package

An exchange package is a 16-word block of data stored in a reserved area of memory that contains the initial parameters for a particular computer program. In addition to initializing the program, these parameters also provide continuity if a program stops and restarts processing from one section of the program to the next.

The exchange package includes the contents of the address (A) and scalar (S) registers. The contents of the intermediate address (B), intermediate scalar (T), vector (V), vector mask (VM), shared B (SB), shared T (ST), and semaphore (SM) registers are not saved in the exchange package. Data in these registers must be stored and replaced as required by the program supervising the object program or by any program that needs this data.

Program Address Register Field

The program address (P) register contents are stored in the program address register field of the exchange package. There are 32 bits in the P register, the lower 2 bits of which are used to select a particular 16-bit parcel of a memory word. The P register is wide enough to address 1 gigaword of memory in C90 mode and 4 Mwords of memory in Y-MP mode.

The address stored in the P register field is the address of the first instruction that issues when the program that corresponds to this exchange package executes.

Instruction Base Address Register Field

The instruction base address (IBA) register holds the base address of the user's instruction area (the location in memory where a program's instruction area begins). The absolute memory address for an instruction fetch sequence is formed by adding the contents of the IBA register to the 30 high-order bits of the contents of the P register.

Instruction Limit Address Register Field

The instruction limit address (ILA) register holds the limit address of the program's memory image area, which is used to determine the highest absolute memory address that can be accessed during an instruction fetch sequence.

The absolute memory address used in an instruction fetch sequence must be an address between the IBA and ILA specified for the program being executed, or a program range error occurs. If the interrupt-on-program-range-error (IPR) mode is set in the exchange package, this error sets the program-range-error (PRE) interrupt flag. Regardless of the state of the IPR mode, a CPU interrupt will occur.

Data Base Address Register Field

The data base address (DBA) register holds the base address of the user's data area (the location in memory where a program's data area begins). Each time an instruction in the program makes a memory reference, the memory address generated by the instruction is added to the DBA to form the absolute memory address.

Data Limit Address Register Field

The data limit address (DLA) register holds the limit address of the user's data area, which is used to determine the highest absolute memory address the program can use for reading or writing data.

Each time an instruction makes a memory reference, the absolute memory address generated is compared to the DLA and the DBA. The absolute memory address must be between the DBA and the DLA, or an operand range error occurs. If the interrupt-on-operand-range-error (IOR) mode is set in the exchange package, this error sets the operand-range-error (ORE) interrupt flag, causing a CPU interrupt.

An instruction that attempts to read from a memory address outside the limits of the DBA and DLA still issues and finishes, but a zero value is transferred from memory. An instruction that attempts to write to a memory address outside these limits issues, but no write operation occurs.

Interrupt Modes Field

There are 16 user-selectable interrupt modes, which allow the programmer to select the conditions under which the active program can be interrupted. These modes are usually selected in the exchange package, and with the exception of IPR, FEX, and FNX, they must be

enabled by setting the EIM (enable interrupt modes) flag. The EIM flag sets automatically on an exchange to non-monitor mode and clears on an exchange back to monitor mode. While in monitor mode, the EIM flag can be set or cleared by instructions 001302 or 001303, respectively.

The interrupt modes are explained briefly in Table 2-1.

Table 2-1. CRAY C90 Series Interrupt Modes

| Mode | Description |
|------|---|
| IRP | Enables an interrupt if a register parity error is detected while reading data from a register. |
| IUM | Enables an interrupt if an uncorrectable memory error is detected while reading data from memory. |
| IFP | Enables an interrupt if a floating-point error occurs. |
| IOR | Enables an interrupt if an operand range error occurs. |
| IPR | Enables the PRE interrupt flag to set if a program range error occurs. A program range error always causes an exchange, regardless of the state of IPR. This mode is not affected by the EIM flag. |
| FEX | Enables the EEX interrupt flag to set if an error exit instruction (000000) issues. Issuing an error exit instruction always causes an exchange, regardless of the state of FEX. This mode is not affected by the EIM flag. |
| IBP | Enables an interrupt if a breakpoint occurs. |
| ICM | Enables an interrupt if a correctable memory error is detected while reading data from memory. |
| IMC | Enables an interrupt if requested by the maintenance control unit (MCU). The MCU for a CRAY C90 series computer system is the MWS-E. |
| IRT | Enables an interrupt if requested by the real-time clock. |
| IIP | Enables an interprocessor interrupt if requested by another CPU. |
| IIO | Enables an I/O interrupt if SIE is set and this CPU is the lowest-numbered CPU with IIO=1 and EIM=1. |
| IPC | Enables an interrupt if requested by the programmable clock. |
| IDL | Enables an interrupt if a deadlock occurs while the program is not in monitor mode. IDL has no effect in monitor mode. |
| IMI | Enables an interrupt if a monitor mode instruction (001ijk; j≠0) issues while the program is not in monitor mode. IMI has no effect in monitor mode. |
| FNX | Enables the NEX interrupt flag to set if a normal exit instruction (004000) issues. Issuing a normal exit instruction always causes an exchange, regardless of the state of FNX. This mode is not affected by the EIM flag. |

Interrupt Flags Field

There are 16 interrupt flags, with one flag corresponding to each of the 16 user-selectable interrupt modes. If a particular interrupt mode (except IPR, FEX, or FNX) is set and enabled and the specified error occurs, the corresponding interrupt flag is set, forcing an exchange. If the error occurs while the appropriate interrupt mode is set but not enabled, the interrupt is held. This condition can occur only while the program is in monitor mode. Enabling the interrupt modes, either by exchanging to user mode or by issuing instruction 001302, enables the held interrupt to be processed, at which time it sets the corresponding interrupt flag and forces an exchange.

All interrupts or held interrupts, except PCI and ICP, are cleared on any exchange. PCI and ICP interrupts are held until they are cleared by instruction 001405 or 001402, respectively.

Two interrupt flags, deadlock (DL) and monitor instruction interrupt (MII), will set only if the corresponding interrupt modes are set and if the program is in non-monitor mode when the error occurs.

The I/O interrupt (IOI) flag sets only if the system I/O interrupts enabled (SIE) flag is set and if the CPU to be interrupted is the lowest-numbered CPU with IIO interrupt mode set and enabled. The SIE flag can be set by any CPU issuing instruction 001600. After any CPU is interrupted by an I/O interrupt, this flag is cleared, disabling all I/O interrupts. The interrupted CPU resets the SIE flag by issuing instruction 001600 after it has serviced the I/O interrupt.

Three errors always cause an exchange, regardless of the status of the EIM flag: a program range error, issuing instruction 000000, or issuing instruction 004000. The interrupt modes specifying these errors (IPR, FEX, and FNX) are used solely to enable setting the corresponding interrupt flags (PRE, EEX, and NEX respectively) should the appropriate error occur. Setting an interrupt flag in these cases makes it easier to determine the source of the error.

The errors that set interrupt flags are explained briefly in Table 2-2.

Table 2-2. CRAY C90 Series Interrupt Flags

| Flag | Description |
|------|---|
| RPE | The register parity error flag sets if the IRP interrupt mode bit is set and enabled and a parity error occurs during a read operation from a B, T, V, SB, or ST register or from an instruction buffer. |
| MEU | The memory error (uncorrectable) flag sets if the IUM interrupt mode bit is set and enabled and an uncorrectable memory error occurs while reading data from memory. |
| FPE | The floating-point error flag sets if the IFP interrupt mode is set and enabled and a floating-point range error occurs in any of the floating-point functional units. |
| ORE | The operand range error flag sets if the IOR interrupt mode bit is set and enabled and a data reference is made outside the address boundaries specified in the DBA and DLA registers. |
| PRE | The program range error flag sets if the IPR interrupt mode bit is set and an instruction fetch is made outside the address boundaries specified in the IBA and ILA registers. A program range error always causes an exchange, regardless of the state of IPR. |
| EEX | The error exit flag sets if the FEX interrupt mode bit is set and an error exit instruction (000000) issues. Issuing an error exit instruction always causes an exchange, regardless of the state of FEX. |
| BPI | The breakpoint interrupt flag sets if the IBP interrupt mode bit is set and enabled and a write reference is made to an address within the breakpoint range. |
| MEC | The memory error (correctable) flag sets if the ICM interrupt mode bit is set and enabled and a correctable memory error occurs while reading data from memory. |
| MCU | The MCU interrupt flag sets if the IMC interrupt mode bit is set and enabled and the MCU interrupt signal becomes active on I/O channel 40. |
| RTI | The real-time interrupt flag sets if the IRT interrupt mode bit is set and enabled and a real-time interrupt request is received. |
| ICP | The interprocessor interrupt flag sets if the IIP interrupt mode bit is set and enabled and another CPU requests an interrupt of this CPU by issuing instruction 0014j1. |
| IOI | The I/O interrupt flag sets if the SIE bit is set and this CPU is the lowest-numbered CPU with IIO interrupt mode set and enabled when a LOSP or VHISP channel completes a transfer. |
| PCI | The programmable clock interrupt flag sets if the IPC interrupt mode bit is set and enabled and the counter in the programmable clock equals 0. |
| DL | The deadlock interrupt flag sets if the IDL interrupt mode bit is set, the program is not in monitor mode, and a deadlock condition occurs because all CPUs in a cluster are holding issue on a test and set instruction. |
| MII | The monitor instruction interrupt flag sets if the IMI interrupt mode bit is set and a monitor mode instruction (001ijk, j ≠ 0) issues while the program is not in monitor mode. |
| NEX | The normal exit flag sets if the FNX interrupt mode bit is set and a normal exit instruction (004000) issues. Issuing a normal exit instruction always causes an exchange, regardless of the state of FNX. |

Status Field

The status field contains 4 bits used to indicate the state of the CPU at the time an exchange occurs. These status bits are set during program execution and are therefore not user selectable. Table 2-3 briefly describes each of the status bits used.

Table 2-3. CRAY C90 Series Status Field Bit Assignments

| Status | Description |
|-----------|---|
| VNU | The vectors not used bit sets if no vector instructions (077 <i>ijk</i> or 140 <i>ijk</i> through 177 <i>ijk</i>) were issued during the execution interval. |
| FPS | The floating-point status bit sets if a floating-point error occurred during the execution interval. |
| WS | The waiting on semaphore bit sets if a test and set instruction (0034 <i>jk</i>) is holding issue in the CIP register. |
| PS or BML | The program state bit is set by the operating system to denote whether a CPU concurrently processing a program with another CPU is the master or slave in a multitasking situation. In CPUs with a BMM functional unit, the PS bit in the status register is used as the B matrix loaded (BML) flag to indicate to the software that the BMM functional unit is loaded. |

Modes Field

There are four user-selectable modes that allow the programmer to select several modes of operation for the program. These modes are described briefly in Table 2-4.

Table 2-4. CRAY C90 Series Operating Modes

| Mode | Description |
|------|---|
| C90 | If C90 mode is set, the program can use the full CRAY C90 series instruction set; otherwise, only CRAY Y-MP instructions can be executed. |
| ESL | If enable second vector logical mode is set, the second vector logical functional unit is enabled, and if it is not busy, it has first priority to execute instructions 140 <i>ijk</i> through 145 <i>ijk</i> . |
| BDM | If bidirectional memory mode is set, block read and write operations can operate concurrently. |
| MM | If monitor mode is set, the program can execute those instructions that are privileged to monitor mode. |

Processor Number Field

The contents of the 4-bit processor number field indicate which CPU performed the exchange sequence. This value is not initially stored in the exchange package before the program starts; it is a constant value inserted into the exchange package after the program has run and been exchanged.

Cluster Number Field

The 5-bit cluster number (CLN) field contains the number to be loaded into the CLN register. This number selects one of 17_{10} available clusters of shared registers that the CPU can access. If the contents of the CLN register are 0, the CPU does not have access to any shared registers. The contents of the CLN registers in all CPUs are also used to determine a deadlock interrupt condition.

Exchange Address Register Field

The 8-bit exchange address (XA) register field specifies the address of the first word of a 16-word exchange package loaded by an exchange sequence. The XA register contains only the 8 high-order bits of a 12-bit absolute memory address. The low-order bits of the address are always 0, because an exchange package must begin on a 16-word boundary. The 12-bit limit on the absolute memory address means that the exchange package area is located in the lower 4,096 (10000_8) words of memory.

Vector Length Register Field

The 8-bit vector length (VL) register field specifies the length of all vector operations performed by vector instructions and the effective number of elements held in the V registers. The value in the VL register can be changed during program execution by using the $00200k$ instruction.

A Register Fields

The current contents of all A registers are stored in bits 2^0 through 2^{31} of words 0 through 7 during an exchange sequence.

S Register Fields

The current contents of all S registers are stored in bits 2^0 through 2^{63} of words 8 through 15 during an exchange sequence.

Instruction Fetch Sequence

An instruction fetch sequence retrieves program code from memory and places it in an instruction buffer. The program code is held in the instruction buffer before being delivered to the instruction issue registers.

Instruction Issue

An instruction issue sequence is the series of steps performed to move an instruction from an instruction buffer through the issue registers and into execution.

Programmable Clock

Each CPU has a programmable clock that generates periodic interrupts at specific preset intervals. Available intervals range between 9 and $2^{32} - 1$ CPs. Intervals shorter than 100 μ s are not practical because of the monitor overhead involved in processing the interrupt. These instructions are privileged to monitor mode.

Status Registers

Each CPU contains eight status registers. Memory error and register parity error information is reported to these registers, as well as information on the status of several bits in the active exchange package.

Performance Monitor

The performance monitor tracks groups of hardware-related events. These results can be used to indicate the relative performance of a program. The performance monitor contains thirty-two 48-bit performance counters.

Performance events are monitored only when operating in non-monitor mode. Entering monitor mode disables the performance counters.

Two types of instructions are used with the performance monitor: user instructions and maintenance instructions. The user instructions allow the user to select and read the performance monitor. The maintenance instructions test the logic of the performance monitor.

Parallel Processing Features

A CRAY C90 series mainframe has several special features that enhance the parallel processing capabilities inherent in the system. Parallel processing can mean different things in different environments; the following subsections discuss two types of parallel processing used:

- Parallel processing within a single CPU
- Parallel processing between two or more CPUs

Parallel processing features within a single CPU include instruction pipelining and segmentation, functional unit independence, vector processing (described earlier in this section), multitasking, and Autotasking. The first two features are inherent hardware features of a CRAY C90 mainframe; a programmer has little control over these features. The vector processing feature can be manipulated by the programmer to provide optimum throughput. Refer to the "Vector Processing" subsection later in this subsection for more information on vector processing.

Parallel processing between two or more CPUs is called multiprocessing: the capability for several programs to run concurrently on multiple CPUs of a single mainframe. Included in this category are multitasking and the Autotasking feature of the CF77 Fortran compiling system. Multitasking is the capability to run two or more parts (or tasks) of a single program in parallel on different CPUs within a mainframe. Autotasking is automatic multiprocessing; it enables user programs to be automatically partitioned over multiple CPUs.

Pipelining and Segmentation

Pipelining means operation or instruction begins before a previous operation or instruction finishes. Pipelining is accomplished using fully segmented hardware. Segmentation means an operation is divided into a discrete number of sequential steps, or segments. Fully segmented hardware is designed to perform one segment of an operation during a single CP. At the beginning of the second CP, the partial results are sent to the second hardware segment in order to process the second step of the operation. During this second CP, the first hardware segment begins to perform the first step of the next operation.

In a CRAY C90 series mainframe, all hardware is fully segmented. Therefore, pipelining occurs during all hardware operations such as exchange sequences, memory references, instruction fetch sequences, instruction issue sequences, and functional unit operations. The pipelining and segmentation features are critical to the execution of vector instructions.

Figure 2-15 shows the pipelining of three sets of scalar instructions through a segmented functional unit.

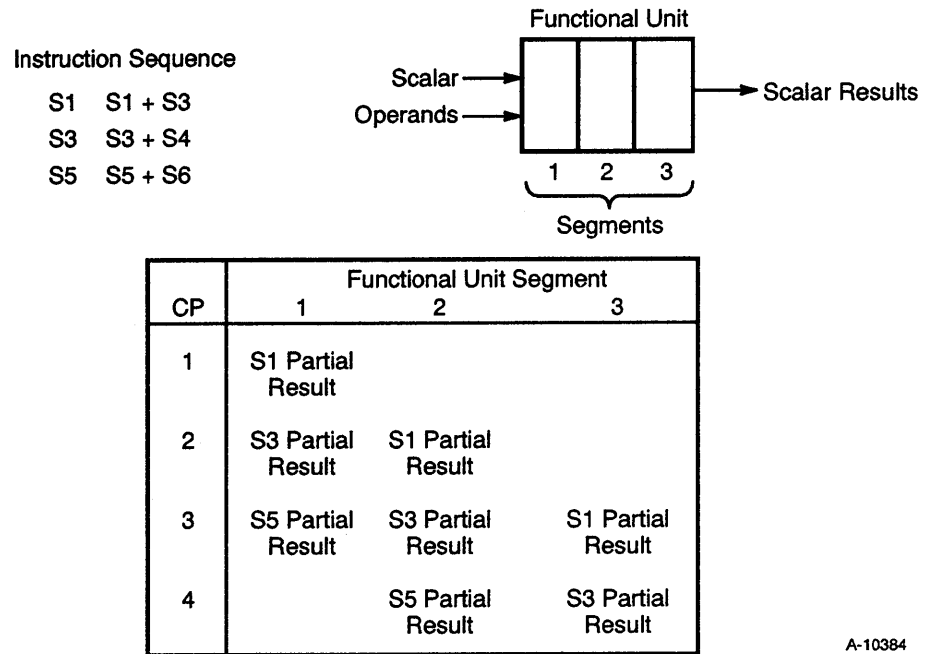
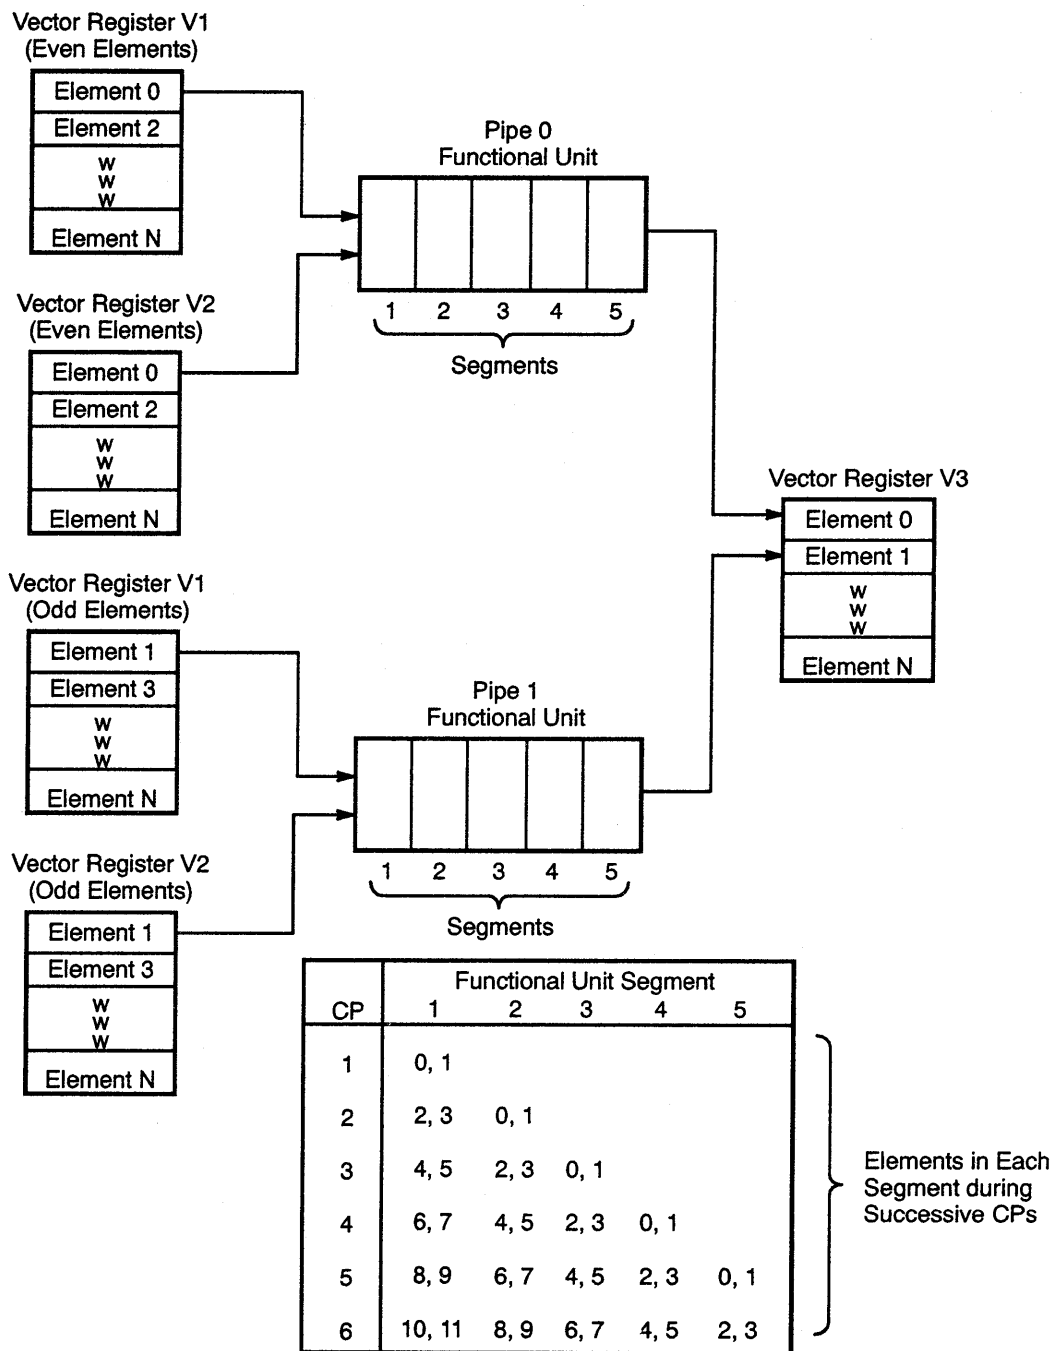


Figure 2-15. Scalar Segmentation and Pipelining Example

In the first CP, the first set of operands enters the first segment of the functional unit. During the next CP, the partial result is moved to the second segment of the functional unit, and the second pair of operands enters the first segment. This process continues each CP until the three operand pairs are completely processed. After 3 CPs, the first result leaves the functional unit and enters scalar register S1; the S3 and S5 results will be available in successive CPs.

A CRAY C90 series mainframe contains two sets of vector functional units: one for processing even-numbered elements and one for processing odd-numbered elements. This enables two pairs of elements to be processed in a single CP and almost doubles the vector processing rate. Figure 2-16 shows how a set of vector elements is pipelined through a dual vector functional unit. In the first CP, element 0 of register V1 and element 0 of register V2 enter the first segment of the pipe 0 functional unit, while element 1 of each register enters the pipe 1 functional unit. During the next CP, the partial results move to the second segments of each functional unit, while element 2 of both vector registers enters the first segment of the pipe 0 functional unit, and element 3 of both vector registers enters the first segment of the pipe 1 functional unit. This process continues each CP until all elements are completely processed.

In this example, the functional units are divided into five segments; the dual functional units process up to ten different pairs of elements simultaneously. After 5 CPs, the first results leave the functional units and enter vector register V3; subsequent results are available at the rate of two results per CP.



A-10385

Figure 2-16. Vector Segmentation and Pipelining Example

Functional Unit Independence

The specialized functional units in a CRAY C90 series mainframe handle the arithmetic, logical, and shift operations. Most functional units are fully independent; any number of functional units can process instructions concurrently. Functional unit independence allows different operations such as multiplications and additions to proceed in parallel.

For example, the operation represented by the equation $A = (B + C) \times D \times E$ could be accomplished as follows. If operands B, C, D, and E are loaded into the S registers, three instructions are generated for the equation: one that adds B and C, one that multiplies D and E, and one that multiplies the results of these two operations. The multiplication of D and E is issued first, followed by the addition of B and C. The addition and multiplication operations proceed concurrently. Because the addition takes less time to run than the multiplication, both operations finish at the same time. The addition operation does not require additional processing time because it occurs during the same time interval as the multiplication operation. The results of these two operations are then multiplied to obtain the final result.

Vector Processing

Vector processing increases processing speed and efficiency by allowing an operation to be performed sequentially on a set (or vector) of operands through the execution of a single instruction.

A vector is an ordered set of elements; each element is represented as a 64-bit word. A vector is distinguished from a scalar, which is a single 64-bit word. Examples of structures in Fortran that can be represented as vectors are one-dimensional arrays and rows, columns, and diagonals of multidimensional arrays. Vector processing occurs when arithmetic or logical operations are applied to vectors; it is distinguished from scalar processing in that it operates on many elements rather than on one.

In vector processing, two successive pairs of elements are processed each CP. The dual vector pipes and the dual sets of vector functional units allow a pair of even-numbered elements and a pair of odd-numbered elements to be processed during the same CP. As each pair of operations is completed, the results are delivered to successive even- or odd-numbered elements of the result register. The vector operation continues until the number of elements processed by the instruction equals the count specified by the vector length (VL) register.

Parallel vector operations allow the generation of more than two results per CP. Parallel vector operations occur automatically in the following situations:

- When successive vector instructions use different functional units and different V registers.
- When successive vector instructions use the result stream from one vector register as the operand of another operation using a different functional unit. This process is known as chaining and is explained later in this subsection.

Advantages of Vector Processing

In general, vector processing is faster and more efficient than scalar processing. Vector processing reduces the overhead associated with maintenance of the loop-control variable (for example, incrementing and checking the count). In many cases, loops processed as vectors are reduced to a simple sequence of instructions without branching backwards. Central memory access conflicts are reduced, and finally, functional unit segmentation is exploited through vector processing because results from the units can then be obtained at the rate of two results per CP.

Vectorization typically speeds up a code segment by an approximate factor of ten. If a segment of code that previously accounted for 50% of a program's running time is vectorized, the overall running time is 55% of the original running time (50% for the unvectorized portion plus $0.1 \times 50\%$ for the vectorized portion). Vectorizing 90% of a program causes running time to drop to 19% of the original execution time.

V Register Functions

The V registers are used solely for vector processing. This is unlike the A and S registers, which are used for many secondary functions. Vector processing allows a single instruction to perform a specified operation sequentially on a set (vector) of operands, to produce a vector of results. Examples of these sets or vectors may be rows or columns of a matrix or elements of a table.

The contents of a V register are transferred to or from central memory by means of a block transfer. A vector block transfer is accomplished by specifying a first word address in central memory, an increment or decrement value for the central memory address, and a vector length. The transfer begins with the first element of the V register and proceeds at a maximum rate of two words per clock period (CP); this rate can be affected by central memory conflicts. A central memory conflict interrupts the vector data stream and can occur in chained operations

(although they do not inhibit chaining). Any interruption in the vector data stream adds proportionally to the total execution time of vector operations.

Single-word data transfers can also be made between an S register and an element of a V register.

Vector Instructions

Vector instructions reference V registers by specifying the register number in the *i*, *j*, or *k* field of the instruction. Refer to the "Instruction Formats" subsection later in this section for information on instruction fields. Operations on vector registers always start with element 0. Individual elements of a V register are designated by octal numbers ranging from 00 through 177. These numbers appear as subscripts to vector register references. For example, V6₂₇ refers to element 27 of V register 6.

Vector instructions reserve V registers as either operands or results. If the register is reserved as an operand, it cannot be used as an operand or result until the operand reservation clears. A vector register can be used as both an operand and result register for the same vector instruction. If a register is reserved as a result, it can be used as an operand through a process called chaining. Refer to the subsection "Vector Chaining" in this section for more information on chaining.

No reservation is placed on the VL register during vector processing. If a vector instruction uses an S register as an operand, no reservation is placed on the S register. Conflicts can occur between vector and scalar operations involving floating-point operations and memory access. With the exception of these operations, the floating-point functional units are always available for scalar operations. The S and VL registers can be modified after the vector instruction issues without affecting the vector operation. The A0 and Ak registers in a vector memory reference can also be modified after the instruction issues.

Because most transfers to or from registers are done in blocks of data, instructions that transfer data between V registers and central memory reserve a port, and functional unit instructions reserve the appropriate functional unit.

Vector Chaining

A CRAY C90 series mainframe allows a vector register reserved for results to become the operand register of a succeeding instruction. This process, called chaining, allows a continuous stream of operands to flow

through the vector registers and functional units. Even when a vector load operation pauses because of memory conflicts, chained operations may proceed as soon as data is available.

This chaining mechanism allows chaining to begin at any point in the result vector data stream. The amount of concurrency in a chained operation depends on the relation between the issue time of the chaining instruction and the arrival time of the result data stream. For full chaining to occur, the chaining instruction must issue and be ready to use element 0 of the result at the same time element 0 arrives at the V register. Partial chaining occurs if the chaining instruction issues after the arrival of element 0 of the result vector data stream.

Figure 2-17 shows how the results of four instructions are chained together. The instruction chaining sequence performs the following operations:

1. Reads a vector of integers from central memory to register V0.
2. Adds the contents of register V0 to the contents of register V1 and sends the results to register V2.
3. Shifts the results obtained in Step 2 and sends the results to register V3.
4. Forms the logical product of the shifted sum obtained in Step 3 with the contents of register V4 and sends the results to register V5.

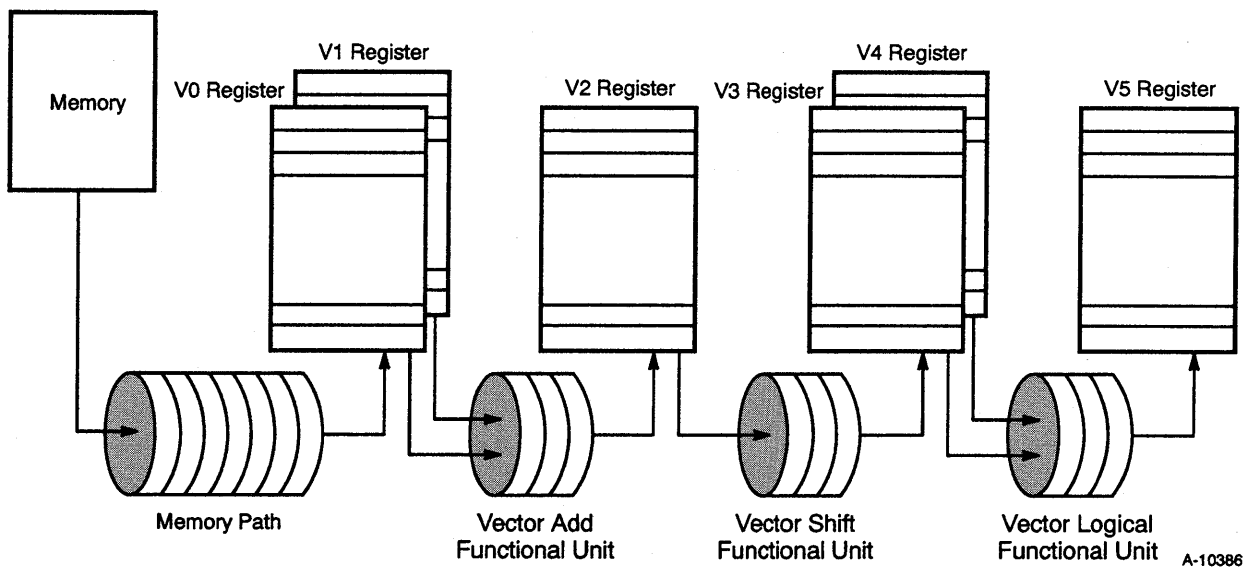


Figure 2-17. Vector Chaining Example

As soon as the first two elements from central memory arrive at register V0, they are added to the first two elements of vector register V1. Subsequent pairs of elements are pipelined through the segmented functional unit, so that a continuous stream of results is sent to the destination register, which is register V2. As soon as the first two elements arrive at register V2, they are used as operands for the shift operation. The results are sent to register V3, which immediately becomes the source of one of the operands necessary for the logical operation between registers V3 and V4. The results of the logical operation are then sent to register V5.

Multiprocessing and Multitasking

Users of CRAY C90 series mainframes can take advantage of parallel processing features known as multiprocessing and multitasking; this category also includes microtasking.

Parallel processing between two or more CPUs is called multiprocessing: the capability for several programs to be run concurrently on multiple CPUs of a single mainframe. Up to n programs can run simultaneously on a machine with n CPUs.

Multitasking is a more recent and complex enhancement than vectorization. Multitasking is the capability to run two or more parts, or tasks, of a single program in parallel on different CPUs within a mainframe. To take advantage of this feature, a program must be logically or functionally divided to allow two or more tasks to run simultaneously (that is, in parallel). For example, a weather modeling application in which the northern hemisphere calculation is one section of code and the southern hemisphere is another section of code. Distinct code segments are not needed; the same code could run on multiple processors simultaneously, with each processor handling separate data. Theoretically, the gain from multitasking can be calculated in the following manner. A program running on a dedicated system in wall clock time (t) could run in a time as short as t/n if multitasked, or modified to use n or more parallel tasks on a machine with n CPUs.

Actually, a speed-up factor of n is not quite attainable because of the additional processing operations (overhead) needed to implement multitasking. In some instances, multitasking can actually increase a program's execution time if the multitasking overhead decreases performance more than parallel processing time improves it. This is a situation that must be investigated before investing too much time and effort.

The following list includes some factors that limit the maximum improvement potential of a program:

- When not all parts of a program can be divided into parallel tasks.
- When those parts that can be multitasked may have dependencies on one another that result in one or more tasks having to wait for other tasks to finish.
- When the multitasking feature incurs additional processing time that is added to the program.

The CFT compiler on a CRAY C90 series mainframe automatically uses the vector hardware to perform operations on inner DO loops that have no data dependencies. Once such optimizing is complete, a single processor can work no faster, but more than one processor could operate on separate parts of the data simultaneously to achieve results faster. Microtasking permits multiple processors to work on a Fortran program at the DO-loop level. The name microtasking was chosen because multiprocessing is efficient even at a DO-loop level where the task size, or granularity, may be small.

Microtasking also works well when the number of processors available is unknown or may vary during program execution. This means that microtasked jobs do not require a dedicated system, although they perform best in a dedicated environment with no competing jobs.

As stated before, advanced programming skills and tools are needed to use multiprocessing, multitasking, and microtasking concepts efficiently.

Autotasking

Analysts and programmers can use Autotasking (automatic multitasking) to automatically detect whether portions of their programs can be run in parallel on a CRAY C90 series mainframe. Autotasking is an extension of multiprocessing and microtasking and is designed to make parallel processing easier to use. Autotasking alters a Fortran program to allow it to run simultaneously on multiple CPUs.

Autotasking is available on CRAY Y-MP computer systems beginning with UNICOS release 4.0 and CF77 release 3.0. Refer to *CF77 Compiling System, Volume 4: Parallel Processing Guide*, Cray Research publication number SG-3074, for more detailed information on Autotasking.


CPU Instructions

The following subsections explain the instruction formats, instruction differences between the Y-MP mode and C90 mode, special register values, special Cray Assembly Language (CAL) syntax forms, and monitor mode instructions used by a CRAY C90 series computer system. A CPU instruction summary is also included.

Notational Conventions

The following conventions are used throughout this section:

- All numbers are decimal numbers unless otherwise indicated.
- Letters X or x or *x* represent an unused value.
- Register bits are numbered from right to left as powers of 2.
- The letter n represents a specified value.
- The notation (value) specifies the contents of a register or memory location as designated by value.
- Variable parameters are in *italic* type.
- The vector mask bits are contained in the VM and VM1 registers. The bits of the VM register correspond to vector elements 0 through 63, and the bits of the VM1 register correspond to vector elements 64 through 127, as shown in Figure 2-18.



A-10387

Figure 2-18. Vector Mask Bits

Instructions can be 1 parcel (16 bits), 2 parcels (32 bits), or 3 parcels (48 bits) long. Instructions are packed 4 parcels per word, and parcels are numbered 0 through 3 from left to right. Any parcel position can be addressed by branch instructions. A 2- or 3-parcel instruction can begin in any parcel of a word and can span a word boundary. For example, a 2-parcel instruction beginning in parcel 3 of word 1 ends in parcel 0 of

word 2. Parcels 0, 1, and 2 of word 1 do not need to be filled with all zeros or ones (padded). Figure 2-19 shows the general instruction format.




Figure 2-19. General Instruction Format

The first parcel is divided into five fields, and the second and third parcels each contain a single field. The four variations of this general format are listed below.

- 1-parcel instruction format with discrete *j* and *k* fields
- 1-parcel instruction format with combined *j* and *k* fields
- 2-parcel instruction format with combined *i*, *j*, *k*, and *m* fields (Y-MP mode only)
- 3-parcel instruction format with combined *m* and *n* fields

Each format uses the fields differently and is described in detail in the following subsections.

Instruction Formats

The following subsections explain the instruction formats, as well as the instruction differences between Y-MP mode and C90 mode.

1-parcel Instruction Format with Discrete *j* and *k* Fields

The most common of the 1-parcel instruction formats uses the *i*, *j*, and *k* fields as individual designators for operand and result registers (refer to Figure 2-20). The *g* and *h* fields define the operation code, the *i* field designates a result register, and the *j* and *k* fields designate operand registers. Some instructions ignore one or more of the *i*, *j*, and *k* fields.




Figure 2-20. 1-parcel Instruction Format with Discrete *j* and *k* Fields

The following types of instructions use this format:

- Arithmetic
- Logical
- Vector shift
- Scalar double shift
- Floating-point constant

1-parcel Instruction Format with Combined *j* and *k* Fields

Some 1-parcel instructions use the *j* and *k* fields as a combined 6-bit field (refer to Figure 2-21). The *g* and *h* fields contain the operation code, and the *i* field usually designates a result register. The combined *j* and *k* fields contain a constant or an intermediate address (B) or intermediate scalar (T) register designator. The 005*ijk* branch instruction and the following types of instructions use the 1-parcel instruction format with combined *j* and *k* fields:

- 6-bit constant
- B or T register block memory transfer
- B or T register data transfer with address (A) or scalar (S) register
- Scalar single shift
- Scalar mask

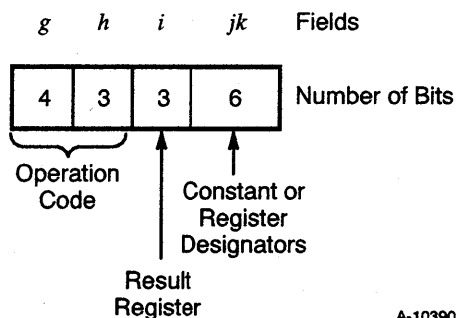


Figure 2-21. 1-parcel Instruction Format with Combined *j* and *k* Fields

2-parcel Instruction Format with Combined *i*, *j*, *k*, and *m* Fields

The 2-parcel instruction format uses the combined *i*, *j*, *k*, and *m* fields to contain a 24-bit address that allows branching to an instruction parcel (refer to Figure 2-22). A 7-bit operation code (*gh*) is followed by an *ijkm* field. The high-order bit of the *i* field (*i*₂) is equal to 0.

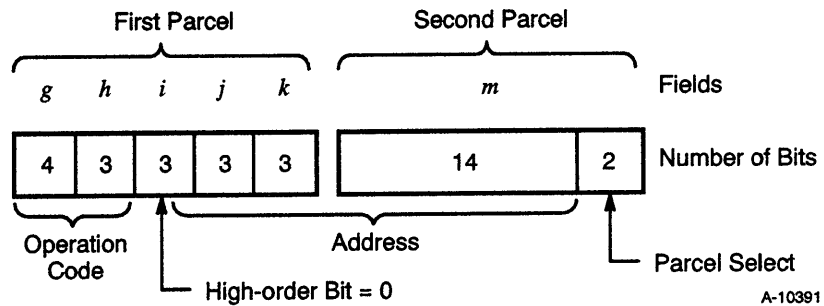


Figure 2-22. 2-parcel Instruction Format with Combined *i*, *j*, *k*, and *m* Fields

3-parcel Instruction Format with Combined *m* and *n* Fields

There are three distinct 3-parcel instruction formats using the combined *m* and *n* fields.

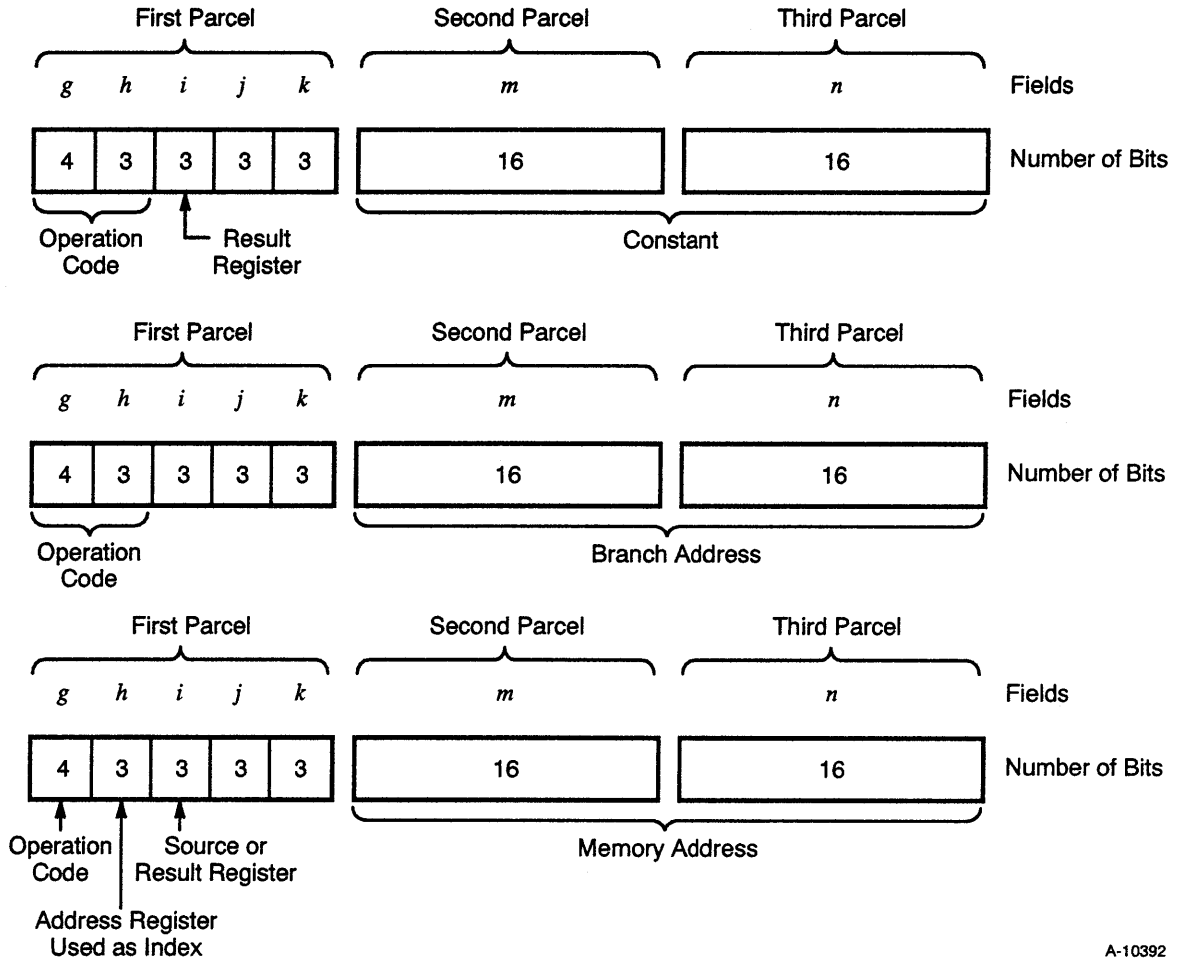
The format for a 32-bit immediate constant uses the combined *m* and *n* fields to hold the constant. The 7-bit *g* and *h* fields contain an operation code, and the 3-bit *i* field designates a result register. The instructions using this format transfer the 32-bit *mn* constant to an A or S register.

NOTE: The *m* field of the 3-parcel instruction contains bits 2^0 through 2^{15} of the expression, while the *n* field contains bits 2^{16} through 2^{31} of the expression. When the instruction is assembled, the *mn* field is reversed and actually appears as the *nm* field when used as an expression.

The format for a C90 mode branch instruction uses the combined *m* and *n* fields to hold the memory branch address. The C90 mode is explained in the next subsection. The 7-bit *g* and *h* fields (and, in one case, bit 2^2 of the *i* field) contain an operation code.

The format for A or S register memory references uses the combined *m* and *n* fields to hold the memory reference address. This format uses the 4-bit *g* field for an operation code, the 3-bit *h* field to designate an address index register, and the 3-bit *i* field to designate a source or result register.

Figure 2-23 shows the three applications for the 3-parcel instruction format with combined m and n fields. Remember that the m and n fields are reversed when a 3-parcel instruction is assembled.



A-10392

Figure 2-23. 3-parcel Instruction Format with Combined m and n Fields

Special Register Values

If the S0 and A0 registers are referenced in the h , j , or k fields of certain instructions, the contents of the respective register are not used; instead, a special operand is generated. This special operand is available regardless of existing A0 or S0 reservations, and in this case the reservation on the register is not checked. This special operand does not alter the actual value of the S0 or A0 register. If register S0 or A0 is referenced in the i field as an operand, the value stored in the register is used. The CAL assembler issues a caution-level error message for A0 or S0 when 0 does not apply to the i field. Table 2-5 lists the special register values.

Table 2-5. Special Register Values

| Field | Operand Value |
|-------------|---------------|
| $Ah, h = 0$ | 0 |
| $Aj, j = 0$ | 0 |
| $Ak, k = 0$ | 1 |
| $Sj, j = 0$ | 0 |
| $Sk, k = 0$ | $2^{63} = 1$ |

Special CAL Syntax Forms

Certain machine instructions can be generated from two or more different CAL instructions. Any of the operations performed by special instructions can be performed by instructions in the basic CAL instruction set. For example, the following CAL instructions generate machine instruction 002000, which enters a 1 into the vector length (VL) register:

```
VL A0
VL 1
```

The first instruction is the basic form of the enter VL instruction, which takes advantage of the special case where $(Ak)=1$ if $k=0$. The second instruction is a special syntax form providing the programmer with a more convenient notation for the special case.

Monitor Mode Instructions

The monitor mode instructions (channel control, set real-time clock, programmable clock interrupts, and so on) perform specialized functions that are useful to the operating system. These instructions run only when the CPU is operating in monitor mode. If a monitor mode instruction issues while the CPU is not in monitor mode, it is treated as a no-operation instruction.

In several cases, a single CAL instruction can generate several different machine instructions. These cases provide for entering the value of an expression into an A register or an S register, or for shifting S register contents. The assembler determines which instruction to generate from characteristics of the expression.

Program Range

The program range, or maximum program length, is 4 Mwords in Y-MP mode and 1 Gword in C90 mode. An instruction outside these ranges produces an undefined result.

CPU Instruction Summary

This subsection introduces and summarizes all mainframe instructions used by a CRAY C90 series computer system. The instructions are summarized two ways: by the functional unit that executes the instruction and by the function the instruction performs.

The following instruction summaries use the acronyms and abbreviations that were defined in previous sections. A glossary is provided at the end of this manual; acronyms and abbreviations are defined there.

In some instructions, register designators are prefixed by the following letters that have special meaning to the assembler. The letters and their meanings are listed as follows:

| <u>Letter</u> | <u>Description</u> |
|---------------|---|
| F | Floating-point operation |
| H | Half-precision floating-point operation |
| I | Reciprocal iteration |
| P | Population count |
| Q | Parity count |
| R | Rounded floating-point operation |
| Z | Leading-zero count |

| <u>Character</u> | <u>Operation</u> |
|------------------|--|
| + | Arithmetic sum of specified registers |
| - | Arithmetic difference of specified registers |
| * | Arithmetic product of specified registers |
| / | Reciprocal approximation |
| # | Use one's complement |
| > | Shift value or form mask from left to right |
| < | Shift value or form mask from right to left |
| & | Logical product of specified registers |
| ! | Logical sum of specified registers |
| \ | Logical exclusive OR of specified registers |

An expression (*exp*) occupies the *jk*, *ijkm*, or *mn* field. The *h*, *i*, *j*, and *k* designators indicate the field of the machine instruction into which the register designator constant or symbol value is placed.

Functional Units Instruction Summary

Instructions other than simple transmit or control operations are performed by specialized hardware components known as functional units. Listed below are the machine instructions performed by each of the functional units.

| <u>Functional Unit</u> | <u>Instructions</u> |
|--------------------------------|----------------------------------|
| Address add (integer) | 030, 031 |
| Address multiply (integer) | 032 |
| Scalar add (integer) | 060, 061 |
| Scalar logical | 042 through 051 |
| Scalar shift | 052 through 057 |
| Scalar pop/parity/leading zero | 026, 027 |
| Vector add (integer) | 154 through 157 |
| Vector logical | 140 through 147, 175 |
| Second vector logical | 140 through 145 |
| Vector shift | 150 through 153 |
| Vector pop/parity | 174ij1, 174ij2 |
| Second vector pop/parity | 174ij1, 174ij2 |
| Bit matrix multiply | 070ij6, 1740j4, 174ij6, 002210 |
| Floating-point add | 062, 063, 170 through 173 |
| Floating-point multiply | 064 through 067, 160 through 167 |
| Floating-point reciprocal | 070, 174ij0 |
| Memory (scalar) | 10h through 13h |
| Memory (vector) | 176, 177 |

Functional Instruction Summary

This subsection summarizes the instructions by the functions they perform. Included is a brief, general description of the function of each group of instructions; then the machine instruction, the CAL syntax, and a description is listed.

NOTE: The following superscripts in the machine instruction column are used throughout the instruction summary:

| <u>Superscript</u> | <u>Description</u> |
|--------------------|----------------------------|
| 1 | C90 mode only |
| 2 | Y-MP mode only |
| 3 | Privileged to monitor mode |

Register Entry Instructions

The register entry instructions transmit values, such as constants, expression values, or masks, directly into registers.

Transfers into A Registers

The following instructions transmit values into the A registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---------------------------------------|
| 020i00 nm | Ai exp | Transmit nm to Ai |
| 021i00 nm | Ai exp | Transmit one's complement of nm to Ai |
| 022ijk | Ai exp | Transmit jk to Ai |
| 031i00 | Ai -1 | Transmit -1 to Ai; (Ai = 77777777) |

Transfers into S Registers

The following instructions transmit values into the S registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 040i00 nm | Si exp | Transmit nm to Si 0 – 31 (32 – 63 = 0) |
| 040i20 nm | Si Si:exp | Transmit nm to Si 00 – 31 (32 – 63 unchanged) |
| 040i40 nm | Si exp:Si | Transmit nm to Si 32 – 63 (00 – 31 unchanged) |
| 041i00 nm | Si exp | Transmit one's complement of nm to Si (32 – 63 = 1) |
| 042i00 | Si -1 | Enter -1 into Si; (Si = 177777 177777 177777 177777) |
| 042ijk | Si <exp | Form ones mask exp bits in Si from the right; jk field contains 100 ₈ -exp |
| 042ijk | Si #>exp | Form zeros mask exp bits in Si from the left; jk field get exp |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 042i77 | Si 1 | Enter 1 into Si |
| 043i00 | Si 0 | Clear Si |
| 043ijk | S > exp | Form ones mask exp bits in Si from the left; jk field gets exp |
| 043ijk | Si #<exp | Form zeros mask exp bits in Si from the right; jk field contains 100g -exp |
| 047i00 | Si #SB | Enter one's complement of sign bit into Si |
| 051i00 | Si SB | Enter sign bit into Si |
| 071i30 | Si 0.6 | Transmit constant 0.75*2**48 to Si (Si = 040060 140000 000000 000000) |
| 071i40 | Si 0.4 | Transmit constant 0.4 to Si (Si = 040000 100000 000000 000000) |
| 071i50 | Si 1.0 | Transmit constant 1.0 to Si (Si = 040001 100000 000000 000000) |
| 071i60 | Si 2.0 | Transmit constant 2.0 to Si (Si = 040002 100000 000000 000000) |
| 071i70 | Si 4.0 | Transmit constant 4.0 to Si (Si = 040003 100000 000000 000000) |

Transfers into V Registers

The following instructions transmit values into the V registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|-----------------------|
| 077i0k | Vi,Ak 0 | Clear Vi element (Ak) |
| 145iii | Vi 0 | Clear Vi |

Transfers into Semaphore Registers

The following instructions transmit values into the semaphore registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 0034jk | SMjk 1,TS | Test and set semaphore <i>jk</i> ($j2 = 0$) |
| 0034jk | SM,Ak 1,TS | Test and set semaphore (<i>Ak</i>) ($j2 = 1$) |
| 0036jk | SMjk 0 | Clear semaphore <i>jk</i> ($j2 = 0$) |
| 0036jk | SM,Ak 0 | Clear semaphore (<i>Ak</i>) ($j2 = 1$) |
| 0037jk | SMjk 1 | Set semaphore <i>jk</i> ($j2 = 0$) |
| 0037jk | SM,Ak 1 | Set semaphore (<i>Ak</i>) ($j2=1$) |

Interregister Transfer Instructions

The interregister transfer instructions transmit the contents of one register to another register. In some cases, the register contents can be complemented, converted to floating-point format, or sign extended as a function of the transfer.

Transfers to A Registers

The following instructions transfer the contents of other registers into the A registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|--------------------|--|
| 023ij0 | <i>Ai Sj</i> | Transmit (<i>Sj</i>) to <i>Ai</i> |
| 023i01 | <i>Ai VL</i> | Transmit VL to <i>Ai</i> [VL = 128 (64)] |
| 024ijk | <i>Ai Bjk</i> | Transmit (<i>Bjk</i>) to <i>Ai</i> |
| 026ij4 | <i>Ai SB,Aj,+1</i> | Transmit (SB) designated by (<i>Aj</i>) to <i>Ai</i> , and increment (SB, <i>Aj</i>) by 1 |
| 026ij5 | <i>Ai SBj,+1</i> | Transmit (SB _{<i>j</i>}) to <i>Ai</i> , and increment (SB _{<i>j</i>}) by 1 |
| 026ij6 | <i>Ai SBj,Aj</i> | Transmit (SB) designated by (<i>Aj</i>) to <i>Ai</i> |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 026ij7 | $A_i SB_j$ | Transmit shared B_j to A_i |
| 030i0k | $A_i Ak$ | Transmit (Ak) to A_i |
| 031i0k | $A_i -Ak$ | Transmit the negative of (Ak) to A_i |
| 033i00 | $A_i CI$ | Transmit I/O interrupting channel number to A_i |
| 033ij0 | $A_i CA, A_j$ | Transmit I/O-current address of channel (A_j) to A_i |
| 033ij1 | $A_i CE, A_j$ | Transmit channel status word (A_j) to A_i |

Transfers to S Registers

The following instructions transmit the contents of other registers into the S registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 047i0k | $S_i \#Sk$ | Transmit one's complement of (Sk) to S_i |
| 051i0k | $S_i Sk$ | Transmit (Sk) to S_i |
| 061i0k | $S_i -Sk$ | Transmit negative of (Sk) to S_i |
| 071i0k | $S_i Ak$ | Transmit (Ak) to S_i with no sign extension |
| 071i1k | $S_i +Ak$ | Transmit (Ak) to S_i with sign extension |
| 071i2k | $S_i *Fak$ | Transmit (Ak) to S_i as unnormalized floating-point number (exponent equals 40060) |
| 072i00 | $S_i RT$ | Transmit (RTC) to S_i |
| 072i02 | $S_i SM$ | Transmit (SM) to S_i |
| 072ij3 | $S_i ST_j$ | Transmit (ST_j) to S_i |
| 073i00 | $S_i VM$ | Transmit vector mask lower to S_i |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 074ijk | $S_i T_{jk}$ | Transmit (T_{jk}) to S_i |
| 076ijk | $S_i V_j, A_k$ | Transmit (V_j , element (A_k)) to S_i |

Transfers to V Registers

The following instructions transmit the contents of other registers into the V registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 077ijk | $V_i, A_k S_j$ | Transmit (S_j) to V_i element (A_k) |
| 142i0k | $V_i V_k$ | Transmit (V_k) to V_i |
| 156i0k | $V_i -V_k$ | Transmit negative of (V_k) to V_i |

Transfers to Intermediate Registers

The following instructions transmit the contents of A and S registers into the intermediate B and T registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--------------------------------|
| 025ijk | $B_{jk} A_i$ | Transmit (A_i) to B_{jk} |
| 075ijk | $T_{jk} S_i$ | Transmit (S_i) to T_{jk} |

Transfers to Shared Registers

The following instructions transmit the contents of other registers into the shared registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 027ij6 | $SB, A_j A_i$ | Transmit (A_i) to SB designated by (A_j) |
| 027ij7 | $SB_j A_i$ | Transmit (A_i) to shared B_j |
| 073i02 | $SM S_i$ | Transmit (S_i) to the semaphore registers |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 073ij3 | STj Si | Transmit (Si) to shared Tj |
| 073ij6 | ST,Aj Si | Transmit (Si) to ST designated by (Aj) |

Transfers to Status Registers

The following instructions transfer the contents of an S register into the status registers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 073i05 | SR0 Si | Transmit (Si) to status register 0 |
| 073i75 ³ | SR7 Si | Transmit (Si) to maintenance mode register (status register 7) |

Transfer to Vector Mask Register

The following instructions transmit the contents of other registers into the vector mask register.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 0030j0 | VM Sj | Transmit (Sj) to vector mask lower $Vm = (2^{127} - 2^{64})$, lower elements 0 – 63 |
| 0030j1 ¹ | VM1 Sj | Transmit (Sj) to vector mask upper $Vm = (2^{63} - 2^0)$, upper elements 64 – 127 |

Transfer to Vector Length Register

The following instructions transmit the contents of other registers into the vector length register.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---------------------|
| 00200k | VL Ak | Transmit (Ak) to VL |
| 002000 | VL 1 | Transmit 1 to VL |

Memory Transfer Instructions

The memory transfer instructions enable or disable bidirectional memory transfers, transfer data between registers and memory, and ensure completion of memory references.

Bidirectional Memory Transfers

The following instructions enable or disable bidirectional memory transfers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 002500 | DBM | Disable bidirectional memory transfers (BDM = 0) |
| 002600 | EBM | Enable bidirectional memory transfers (BDM = 1) |

Memory References

The following instructions ensure completion of instructions for bidirectional memory transfers.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--------------------------------|
| 002700 | CMR | Complete memory references |
| 002704 | CPA | Complete port reads and writes |
| 002705 | CPR | Complete port reads |
| 002706 | CPW | Complete port writes |

Writes

The following instructions write values into memory.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 035ijk | ,A0 Bjk,Ai | Write (Ai) words from B register <i>jk</i> to memory address ((A0) + (DBA)) |
| 037ijk | ,A0 Tjk,Ai | Write (Ai) words from T register <i>jk</i> to memory address ((A0) + (DBA)) |
| 11hi00 nm | exp,Ah Ai | Write (Ai) to memory address ((Ah) + exp + (DBA)) |
| 13hi00 nm | exp,Ah Si | Write (Si) to memory address ((Ah) + exp + (DBA)) |
| 1770jk | ,A0,Ak Vj | Write (VL) words from Vj to memory address ((A0) + (DBA)) incremented by (Ak) |
| 1770j0 | ,A0, 1 Vj | Write (VL) words from Vj to memory address ((A0) + (DBA)) incremented by 1 |
| 1771jk | ,A0,Vk Vj | Write (VL) words from Vj to memory address ((A0) + (Vk) + (DBA)) (scatter) |

Reads

The following instructions read values from memory.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 034ijk | Bjk,Ai ,A0 | Read (Ai) words to B register <i>jk</i> from memory address ((A0) + (DBA)) Port A |
| 036ijk | Tjk,Ai ,A0 | Read (Ai) words to T register <i>jk</i> from memory address ((A0) + (DBA)) Port B |
| 10hi00 nm | Ai exp,Ah | Read from memory ((Ah) + exp + (DBA)) to Ai |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 12hi00 nm | Si exp,Ah | Read from memory address ((Ah) + exp + (DBA)) to Si |
| 176i0k | Vi ,A0,Ak | Read (VL) words to Vi from memory address ((A0) + (DBA)) incremented by (Ak) |
| 176i00 | Vi ,A0,1 | Read (VL) words to Vi from memory address ((A0 + (DBA)) incremented by 1 |
| 176i1k | Vi ,A0,Vk | Read (VL) words to Vi from memory address ((A0 + (Vk) + (DBA)) (gather) |

Integer Arithmetic Instructions

Integer arithmetic operations obtain operands from registers and return results to registers. No direct memory references are allowed.

The assembler recognizes several special syntax forms for increasing or decreasing register contents, such as the operands A_{i+1} and A_{i-1} ; however, these references actually result in register references such that the 1 becomes a reference to A_k with $k = 0$.

All integer arithmetic is two's complement and is represented as such in the registers. The address add and address multiply functional units perform 32-bit arithmetic. The scalar add functional unit and the vector add functional unit perform 64-bit arithmetic. No overflow conditions are detected by functional units when performing integer arithmetic.

Multiplication of two fractional operands is accomplished using a floating-point multiply instruction. The floating-point multiply functional unit recognizes conditions in which both operands have zero exponents as a special case and returns the high-order 48 bits of the result as an unnormalized fraction. Division of integers requires that they first be converted to floating-point format and then divided using the floating-point functional units.

32-bit Integer Arithmetic

The following instructions perform 32-bit integer arithmetic.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 030ijk | $A_i \ A_j + A_k$ | Integer sum of (A_j) and (A_k) to A_i |
| 030ij0 | $A_i \ A_j + 1$ | Integer sum of (A_j) and 1 to A_i |
| 031ijk | $A_i \ A_j - A_k$ | Integer difference of (A_j) and (A_k) to A_i |
| 031ij0 | $A_i \ A_j - 1$ | Integer difference of (A_j) and 1 to A_i |
| 032ijk | $A_i \ A_j * A_k$ | Integer product of (A_j) and (A_k) to A_i |

64-bit Integer Arithmetic

The following instructions perform 64-bit integer arithmetic.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 060ijk | $S_i \ S_j + S_k$ | Integer sum of (S_j) and (S_k) to S_i |
| 061ijk | $S_i \ S_j - S_k$ | Integer difference of (S_j) and (S_k) to S_i |
| 154ijk | $V_i \ S_j + V_k$ | Integer sums of (S_j) and (V_k) to V_i |
| 155ijk | $V_i \ V_j + V_k$ | Integer sums of (V_j) and (V_k) to V_i |
| 156ijk | $V_i \ S_j - V_k$ | Integer differences of (S_j) and (V_k) to V_i |
| 157ijk | $V_i \ V_j - V_k$ | Integer differences of (V_j) and (V_k) to V_i |

Bit Matrix Multiply

The following instructions perform bit matrix multiply operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|------------------------|---|
| 174xj4 | B Vj | Load (VL) elements of (Vj) as the transpose of matrix B |
| 174xj6 | Vi Vj * B ^t | Logical bit matrix multiply of (VL) elements of (Vj) and B ^t into Vi |
| 070ij6 | Si Sj * B ^t | Logical bit matrix multiply of (Sj) and B ^t into Si |
| 002210 | CBL | Clear bit matrix loaded flag |

Floating-point Arithmetic Instructions

All floating-point arithmetic operations use registers as the source of operands and return results to registers.

Floating-point numbers are represented in a standard format throughout the CPU. This format is a packed representation of a binary coefficient and an exponent or power of 2. The coefficient is a 48-bit signed fraction. The sign of the coefficient is separated from the rest of the coefficient. Because the coefficient is signed magnitude, it is not complemented for negative values. Refer to "Floating-point Arithmetic" earlier in this section for more information on floating-point numbers and arithmetic.

Floating-point Range Errors

The following instructions enable or disable floating-point range errors to be flagged.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 002100 | EFI | Enable FP Interrupt (IFP = 1, Clear FPS) |
| 002200 | DFI | Disable FP Interrupt (IFP = 0, Clear FPS) |

Floating-point Addition and Subtraction

The following instructions perform floating-point addition or subtraction.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 062ijk | $S_i S_j + FS_k$ | Floating-point sum of (S_j) and (S_k) to S_i |
| 062i0k | $S_i +FS_k$ | Normalize (S_k) to S_i |
| 063ijk | $S_i S_j - FS_k$ | Floating-point difference of (S_j) and (S_k) to S_i |
| 063i0k | $S_i -FS_k$ | Transmit normalized negative of (S_k) to S_i |
| 170ijk | $V_i S_j + FV_k$ | Floating-point sums of (S_j) and (V_k) to V_i |
| 170i0k | $V_i +FV_k$ | Normalize (V_k) to V_i |
| 171ijk | $V_i V_j + FV_k$ | Floating-point sums of (V_j) and (V_k) to V_i |
| 172ijk | $V_i S_j - FV_k$ | Floating-point differences of (S_j) and (V_k) to V_i |
| 172i0k | $V_i -FV_k$ | Transmit normalized negatives of (V_k) to V_i |
| 173ijk | $V_i V_j - FV_k$ | Floating-point differences of (V_j) and (V_k) to V_i |

Floating-point Multiplication

The following instructions perform floating-point multiplication.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 064ijk | $S_i S_j * FS_k$ | Floating-point product of (S_j) and (S_k) to S_i |
| 065ijk | $S_i S_j * HS_k$ | Half-precision rounded floating-point product of (S_j) and (S_k) to S_i |
| 066ijk | $S_i S_j * RS_k$ | Full-precision rounded floating-point product of (S_j) and (S_k) to S_i |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 160ijk | $V_i S_j * FV_k$ | Floating-point products of (S_j) and (V_k) to V_i |
| 161ijk | $V_i V_j * FV_k$ | Floating-point products of (V_j) and (V_k) to V_i |
| 162ijk | $V_i S_j * HV_k$ | Half-precision rounded floating-point products of (S_j) and (V_k) to V_i |
| 163ijk | $V_i V_j * HV_k$ | Half-precision rounded floating-point products of (V_j) and (V_k) to V_i |
| 164ijk | $V_i S_j * RV_k$ | Rounded floating-point products of (S_j) and (V_k) to V_i |
| 165ijk | $V_i V_j * RV_k$ | Rounded floating-point products of (V_j) and (V_k) to V_i |

Reciprocal Approximation

The following instructions perform floating-point reciprocal approximation operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 070ij0 | S_i /HS_j | Floating-point reciprocal approximation of (S_j) to S_i |
| 174ij0 | V_i /HV_j | Floating-point reciprocal approximations of (V_j) to V_i |

Logical Operation Instructions

The scalar and vector logical functional units perform bit-by-bit manipulation of 64-bit quantities. Logical operations include logical products, logical sums, logical exclusive ORs, logical equivalence, vector mask, and merges. Logical operations are defined below.

- A logical product (& operator) is the AND function.
- A logical exclusive or (\ operator) is the exclusive OR function.

- A logical sum (! operator) is the inclusive OR function.
- A logical merge combines two operands depending on a ones mask in a third operand. The result is defined by (operand 2 & mask) ! (operand 1 & #mask).

Logical Products

The following instructions perform logical product operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|--------------------|--|
| 044ijk | $S_i S_j \& S_k$ | Logical product of (Sj) and (Sk) to Si |
| 044ij0 | $S_i S_j \& SB$ | Sign bit of (Sj) to Si |
| 044ij0 | $S_i SB \& S_j$ | Sign bit of (Sj) to Si (j p 0) |
| 045ijk | $S_i \#S_k \& S_j$ | Logical product of (Sj) and one's complement of (Sk) to Si |
| 045ij0 | $S_i \#SB \& S_j$ | (Sj) with sign bit cleared to Si |
| 140ijk | $V_i S_j \& V_k$ | Logical products of (Sj) and (Vk) to Vi |
| 141ijk | $V_i V_j \& V_k$ | Logical products of (Vj) and (Vk) to Vi |

Logical Sums

The following instructions perform logical sum operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 051ijk | $S_i S_j ! S_k$ | Logical sum of (Sj) and (Sk) to Si |
| 051ij0 | $S_i S_j ! SB$ | Logical sum of (Sj) and sign bit to Si |
| 051ij0 | $S_i SB ! S_j$ | Logical sum of (Sj) and sign bit to Si (j ≠ 0) |
| 142ijk | $V_i S_j ! V_k$ | Logical sums of (Sj) and (Vk) to Vi |
| 143ijk | $V_i V_j ! V_k$ | Logical sums of (Vj) and (Vk) to Vi |

Logical Exclusive ORs

The following instructions perform exclusive OR operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|--------------------------------|--------------------------|---|
| 046ijk | $S_i S_j \backslash S_k$ | Exclusive OR of (S_j) and (S_k) to S_i |
| 046ij0 | $S_i S_j \backslash SB$ | Toggle sign bit of (S_j), then enter into S_i |
| 046ij0 | $S_i SB \backslash S_j$ | Toggle sign bit of (S_j), then enter into S_i ($j \neq 0$) |
| 144ijk | $V_i S_j \backslash V_k$ | Exclusive ORs of (S_j) and (V_k) to V_i |
| 145ijk | $V_i V_j \backslash V_k$ | Exclusive ORs of (V_j) and (V_k) to V_i |

Logical Equivalence

The following instructions perform logical equivalence operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|--------------------------------|----------------------------|--|
| 047ijk | $S_i \#S_j \backslash S_k$ | Logical equivalence of (S_k) and (S_j) to S_i |
| 047ij0 | $S_i \#SB \backslash S_j$ | Logical equivalence of (S_j) and sign bit to S_i ($j \neq 0$) |
| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
| 047ij0 | $S_i \#S_j \backslash SB$ | Logical equivalence of (S_j) and sign bit to S_i |

Vector Mask

The following instructions test the elements of a vector register and use the test results to set the corresponding bits of the vector mask register.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 1750j0 | VM Vj,Z | Set VM bit if (Vj) = 0 |
| 1750j1 | VM Vj,N | Set VM bit if (Vj) ≠ 0 |
| 1750j2 | VM Vj,P | Set VM bit if (Vj) ≥ 0 |
| 1750j3 | VM Vj,M | Set VM bit if (Vj) < 0 |
| 175ij4 | Vi,VM Vj,Z | Set VM bit if (Vj) = 0; also, store the compressed indices of the Vj elements = 0 in the Vi elements |
| 175ij5 | Vi,VM Vj,N | Set VM bit if (Vj) ≠ 0; also, store the compressed indices of the Vj elements ≠ 0 in the Vi elements |
| 175ij6 | Vi,VM Vj,P | Set VM bit if (Vj) ≥ 0; also, store the compressed indices of the Vj elements ≥ 0 in the Vi elements |
| 175ij7 | Vi,VM Vj,M | Set VM bit if (Vj) < 0; also store the compressed indices of the Vj elements < 0 in the Vi elements |

Merge

The following instructions perform a logical merge that combines two operands according to the bits set in a ones mask in a third operand.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 050ijk | Si Sj!Si&Sk | Logical product of (Si) and (Sk) complement ORed with logical product of (Sj) and (Sk) to Si |
| 050ij0 | Si Sj!Si&SB | Scalar merge of (Si) and sign of (Sj) to Si |
| 146ijk | Vi Sj!Vk&VM | Transmit (Sj) if VM bit = 1 or (Vk) if VM bit = 0 to Vi |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 146i0k | $V_i \#VM\&Vk$ | Vector merge of (Vk) and 0 to V_i |
| 147ijk | $V_i V_j!Vk\&VM$ | Transmit (V_j) if VM bit = 1 or (V_k) if VM bit = 0 to V_i |

Shift Instructions

The scalar shift functional unit and vector shift functional unit shift 64-bit quantities or 128-bit quantities. A 128-bit quantity is formed by concatenating two 64-bit quantities. The number of bits a value is shifted left or right is determined by the value of an expression for some instructions and by the contents of an A register for other instructions. If the count is specified by an expression, the value of the expression must not exceed 64.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|-----------------------------|----------------------|---|
| 052ijk | $S_0 S_i < exp$ | Shift (S_i) left $exp = jk$ places to S_0 |
| 053ijk | $S_0 S_i > exp$ | Shift (S_i) right $exp = 100_8 - jk$ places to S_0 |
| 054ijk | $S_i S_i < exp$ | Shift (S_i) left $exp = jk$ places to S_i |
| 055ijk | $S_i S_i > exp$ | Shift (S_i) right $exp = 100_8 - jk$ places to S_i |
| 056ijk | $S_i S_i, S_j < A_k$ | Shift (S_i) and (S_j) left (A_k) places to S_i |
| 056ij0 | $S_i S_i, S_j < 1$ | Shift (S_i) and (S_j) left one place to S_i |
| 056i0k | $S_i S_i < A_k$ | Shift (S_i) left (A_k) places to S_i |
| 057ijk | $S_i S_j, S_i > A_k$ | Shift (S_j) and (S_i) right (A_k) places to S_i |
| 057ij0 | $S_i S_j, S_i > 1$ | Shift (S_j) and (S_i) right one place to S_i |
| 057i0k | $S_i S_i > A_k$ | Shift (S_i) right (A_k) places to S_i |
| 150ijk | $V_i V_j < A_k$ | Shift (V_j) left (A_k) places to V_i |
| 150ij0 | $V_i V_j < 1$ | Shift (V_j) left one place to V_i |
| 005400, 150ij0 ¹ | $V_i V_j < V_0$ | Shift (V_j) left (V_0) places to V_i |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|-----------------------------|----------------------|--|
| 151ijk | $V_i V_j > A_k$ | Shift (Vj) right (Ak) places to Vi |
| 151ij0 | $V_i V_j > 1$ | Shift (Vj) right one place to Vi |
| 005400, 151ij0 ¹ | $V_i V_j > V_0$ | Shift (Vj) right (V0) places to Vi |
| 152ijk | $V_i V_j, V_j < A_k$ | Double shift (Vj) left (Ak) places to Vi |
| 152ij0 | $V_i V_j, V_j < 1$ | Double shift (Vj) left one place to Vi |
| 005400, 152ijk ¹ | $V_i V_j, A_k$ | Vector word shift of (Vj) starting at element (Ak) to Vi ((Ak) < VL) |
| 153ijk | $V_i V_j, V_j > A_k$ | Double shift (Vj) right (Ak) places to Vi |
| 153ij0 | $V_i V_j, V_j > 1$ | Double shift (Vj) right one place to Vi |

Bit Count Instructions

Bit count instructions count the number of set bits or the number of leading 0 bits in an S or V register.

Scalar Population Count

The following instruction performs the scalar population count.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--------------------------------|
| 026ij0 | $A_i P S_j$ | Population count of (Sj) to Ai |

Vector Population Count

The following instruction performs the vector population count.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---------------------------------|
| 174ij1 | $V_i P V_j$ | Population counts of (Vj) to Vi |

Population Parity Count

The following instructions perform population parity count.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|--------------------------------|-------------------|--|
| 026ij1 | A_i QSj | Population count parity of (Sj) to A_i |
| 174ij2 | V_i QVj | Population count parities of (Vj) to V_i |

Scalar Leading Zero Count

The following instruction performs the scalar leading zero count.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|--------------------------------|-------------------|-------------------------------------|
| 027ij0 | A_i ZSj | Leading zero count of (Sj) to A_i |

Vector Leading Zero Count

The following instruction performs the vector leading zero count.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|--------------------------------|-------------------|-------------------------------------|
| 174ij3 | V_i ZVj | Leading zero count of (Vj) to V_i |

Branch Instructions

Instructions in this category include conditional and unconditional branch instructions. An expression or the contents of a B register specify the branch address. An address is always taken to be a parcel address when the instruction runs. If an expression has a word-address attribute, the assembler issues an error message.

Unconditional Branch Instructions

The following instructions perform unconditional branch operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 0050jk | J Bjk | Jump to (Bjk) |
| 0051jk | Jinv Bjk | Jump to (Bjk) (Maintenance only: invalidates instruction buffers) |
| 006ijkm ² | J exp | Jump to exp |
| 006000 nm ¹ | J exp | Jump to exp |

Conditional Branch Instructions

The following instructions perform conditional branch operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 0064jk nm ¹ | JTSjk exp | Branch to exp if (SMjk) = 1; else set SMjk (j2 = 0) |
| 0064jk nm ¹ | JTS,Ak exp | Branch to exp if (SM,(Ak)) = 1; else set SM,(Ak) (j2 = 1) |
| 010ijkm ² | JAZ exp | Jump to exp if (A0) = 0 (i2 = 0) |
| 010000 nm ¹ | JAZ exp | Jump to exp if (A0) = 0 |
| 011ijkm ² | JAN exp | Jump to exp if (A0) ≠ 0 (i2 = 0) |
| 011000 nm ¹ | JAN exp | Jump to exp if (A0) ≠ 0 |
| 012ijkm ² | JAP exp | Jump to exp if (A0) is positive; (A0) ≥ 0 (i2 = 0) |
| 012000 nm ¹ | JAP exp | Jump to exp if (A0) is positive; (A0) ≥ 0 |
| 013ijkm ² | JAM exp | Jump to exp if (A0) is negative (i2 = 0) |
| 013000 nm ¹ | JAM exp | Jump to exp if (A0) is negative |
| 014ijkm ² | JSZ exp | Jump to exp if (S0) = 0 (i2 = 0) |
| 014000 nm ¹ | JSZ exp | Jump to exp if (S0) = 0 |

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 015ijkm ² | JSN <i>exp</i> | Jump to <i>exp</i> if (S0) ≠ 0 (<i>i2</i> = 0) |
| 015000 nm ¹ | JSN <i>exp</i> | Jump to <i>exp</i> if (S0) ≠ 0 |
| 016ijkm ² | JSP <i>exp</i> | Jump to <i>exp</i> if (S0) is positive; (S0) ≥ 0 (<i>i2</i> = 0) |
| 016000 nm ¹ | JSP <i>exp</i> | Jump to <i>exp</i> if (S0) is positive; (S0) ≥ 0 |
| 017ijkm ² | JSM <i>exp</i> | Jump to <i>exp</i> if (S0) is negative (<i>i2</i> = 0) |
| 017000 nm ¹ | JSM <i>exp</i> | Jump to <i>exp</i> if (S0) is negative |

Return Jump

The following instructions perform a return jump operation.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 007ijkm ² | R <i>exp</i> | Return jump to <i>exp</i> and set register B00 to (P) + 2 |
| 007000 nm ¹ | R <i>exp</i> | Return jump to <i>exp</i> and set register B00 to (P) + 3 |

Normal Exit

The following instruction performs a normal exit operation.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--------------------|
| 004000 | EX | Normal Exit |

Error Exit

The following instruction performs an error exit operation.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--------------------|
| 000000 | ERR | Error exit |

Monitor Mode Instructions

Monitor mode instructions are executed only when the CPU is in monitor mode. An attempt to execute one of these instructions when not in monitor mode is treated as a pass instruction. The instructions perform specialized functions useful to the operating system.

Channel Control

The following instructions perform channel control operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---|
| 0010jk ³ | CA,Aj Ak | Set channel (Aj) CA register to (Ak) and begin I/O sequence |
| 001000 ³ | NOP | Pass |
| 0011jk ³ | CL,Aj | Set channel (Aj) CL register to (Ak) |
| 0012j0 ³ | CI,Aj | Clear channel (Aj) interrupt and error flags; clear device master clear (output channel) |
| 0012j1 ³ | MC,Aj | Clear channel (Aj) interrupt and error flags, set device master clear (output channel); clear device ready-held (input channel) |

Set Exchange Address

The following instruction sets the exchange address.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|-----------------------------|
| 0013j0 ³ | XA Aj | Enter XA register with (Aj) |

Set Real-time Clock

The following instruction performs a real-time clock operation.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|------------------------------|
| 0014j0 ³ | RT Sj | Enter RTC register with (Sj) |

Set Cluster Number

The following instruction sets the cluster number.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|---------------------------------|
| 0014j3 ³ | CLN Aj | Transmit (Aj) to cluster number |

Programmable Clock Interrupt

The following instructions perform programmable clock operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 0014j4 ³ | PCI Sj | Transmit (Sj) to programmable clock |
| 001405 ³ | CCI | Clear programmable clock interrupt (PCI) request |
| 001406 ³ | ECI | Enable PCI (MM IPC only) |
| 001407 ³ | DCI | Disable PCI (MM IPC only) |

Operand Range Error Interrupt

The following instructions enable or disable operand range error interrupts.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|-----------------------------------|
| 002300 | ERI | Enable range interrupt (IOR = 1) |
| 002400 | DRI | Disable range interrupt (IOR = 0) |

Interprocessor Interrupt

The following instructions perform interprocessor interrupt operations.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 0014j1 ³ | SIPI Aj | Set interprocessor interrupt of CPU (Aj) |
| 001401 ³ | SIPI | Send interprocessor interrupt to CPU0 |
| 001402 ³ | CIPI | Clear interprocessor interrupt |

Breakpoint Interrupt

The following instructions enable or disable breakpoint interrupts.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 002301 | EBP | Enable breakpoint interrupt (IBP = 1) |
| 002401 | DBP | Disable breakpoint interrupt (IBP = 0) |

Performance Counters

The following instructions operate the performance monitor.

| <u>Machine Instruction</u> | <u>CAL Syntax</u> | <u>Description</u> |
|----------------------------|-------------------|--|
| 001500 ³ | | Clear all performance monitor counters |
| 073i21 ³ | Si SR2 | Read PM counters 00 – 17 and increment pointer |
| 073i25 ³ | SR2 Si | Issue PM maintenance advance |
| 073i31 ³ | Si SR3 | Read PM counters 20 – 37 and increment pointer |

CRAY C90 Series Mainframe Specifications

System Clock

Speed 4.2 ns

Optional Functional Units:

Second vector population/parity/
leading zero 2
Bit matrix multiply 2

CPU Specifications

Number of CPUs † 1 to 16

Number of registers per CPU:

Address (A) registers 8
32 bits each
Intermediate address (B) registers 64
32 bits each
Scalar (S) registers 8
64 bits each
Intermediate scalar (T) registers 64
64 bits each
Vector (V) registers 8
64 bits x 128 elements
(C90 mode)
64 bits x 64 elements
(Y-MP mode)
Vector length (VL) register 1
8 bits (C90 mode)
7 bits (Y-MP mode)
Vector mask (VM, VM1) registers 2
64-bits each
Program address (P) register 1
32 bits (C90 mode)
24 bits (Y-MP mode)

Number of functional units per CPU:

Address addition 1
Address multiplication 1
Scalar addition 1
Scalar shift 1
Scalar logical 1
Scalar population/parity/leading zero .. 1
Vector addition 2
Vector shift 2
Vector population/parity/leading zero .. 2
Full vector logical 2
2nd vector logical 2
Vector population/parity/leading zero .. 2
Floating-point addition 2
Floating-point multiplication 2
Floating-point reciprocal approx. 2

CPU Shared Resources

Input/output section:

Very-high speed channels † 0 to 8
Operation half duplex
Channel width 128 bits
Transfer rate 1,800 Mbytes/s
Data protection SECDED
High-speed channel pairs † 2 to 16
Operation full duplex
Channel width 64 bits
Transfer rate 200 Mbytes/s
Data protection SECDED
Low-speed channel pairs † 2 to 16
Operation full duplex
Channel width 16 bits
Transfer rate 6 Mbytes/s
Data protection parity

Central memory:

Word width 64 bits
SBCDBD error correction 16 bits
Memory size (in Mwords) † . 64 to 1,024
Number of banks † 64 to 1,024
Number of modules † 2 to 16
Number of ports per CPU 4

Number of shared register clusters:

16 CPUs 17
8 CPUs 9
4 CPUs 5
2 CPUs 3

Number of shared registers contained in each cluster:

Shared address (SB) registers 8
32 bits each
Shared scalar (ST) registers 8
64 bits each
Semaphore (SM) registers 32
1 bit each

Real-time clock (64-bits) 1

† This information varies, depending on the system configuration. Refer to "System Configurations" in Section 1 for specific numbers for each model.

3 I/O SUBSYSTEM

The Cray Research input/output subsystem model E (IOS-E) controls data transfers between several components of a CRAY C90 series computer system. The IOS-E transfers data to and receives data from the following components.

- The CRAY C90 series mainframe
- The SSD solid-state storage device model E (SSD-E)
- Peripheral devices such as disk drives and front-end computers
- The operator workstation model E (OWS-E)
- The maintenance workstation model E (MWS-E)

The IOS-E comprises a maximum of eight clusters and two workstation interfaces (WINs). The following subsections describe the I/O clusters and WINs. A block diagram of an I/O cluster is provided in Figure 3-1.

I/O Cluster

Each I/O cluster contains the following components:

- 1 I/O processor multiplexer (IOP MUX)
- 4 auxiliary I/O processors (EIOPs)
- 16 I/O buffers
- 1 low-speed (LOSP) channel pair
- 2 high-speed (HISP) channel pairs
- 16 channel adapters

The IOPs (IOP MUX and EIOPs) provide internal control for the I/O cluster. The LOSP channels allow the I/O cluster to exchange control information with the mainframe. The I/O buffers, HISP channels, and channel adapters provide the paths for data transfers between the IOS-E and the mainframe, SSD-E, and peripheral devices.

The IOPs are identical; they have the same architecture and execute the same instruction set. Each IOP is a high-speed 16-bit (1 parcel) computer designed to efficiently control data transfers. Each IOP contains a 64-Kparcel local memory that is protected with SECDED (single-error correction/double-error detection) logic. Each IOP also contains a 128-parcel operand register file that is parity protected and three programmable registers. Each IOP contains 29 I/O channels; 5 of the channels monitor and control operations within the IOP, and 24 channels enable the IOP to communicate with external devices.

Each IOP executes a set of 128 1-parcel or 2-parcel instructions. Ninety-six instructions perform basic operations such as data transfers; arithmetic (two's complement), logical, and shift operations; conditional and unconditional jumps; and subroutine calls and exits. Thirty-two instructions, called I/O functions, control and monitor the I/O channels.

I/O Buffers

The 16 I/O buffers provide temporary storage for data transferred to or from the IOS-E. Each buffer contains 64 Kwords. Each word is 64 bits long and is protected with SECDED logic. Each buffer can simultaneously pass data to or from a channel adapter while passing data to or from the mainframe or SSD-E. Each buffer is dedicated to one peripheral device, or in the case of mass storage devices, to one group of identical devices.

Each EIOP is dedicated to 4 of the 16 I/O buffers. Each EIOP controls all transfers between its buffers and peripheral devices. The EIOPs work with the IOP MUX to control transfers between the buffer and the mainframe or SSD-E. Each EIOP can also transfer data between its I/O buffers and its local memory.

Low-speed and High-speed Channels

The LOSP and HISP channel pairs enable the I/O cluster to communicate with the mainframe and SSD. The LOSP channel pair transfers control information between the IOP MUX and the mainframe. One HISP channel pair transfers data between the I/O buffers and the mainframe; the second pair transfers data between the I/O buffers and the SSD.

The LOSP channel pair operates at 6 or 20 Mbytes/s. The LOSP channels are 16 bits wide and contain 4 parity bits for error detection. Each channel can operate simultaneously.

The HISP channel pairs operate at 200 Mbytes/s for the SSD-E and 158 Mbytes/s for the SSD-E/32i. The HISP channels are 64 bits wide and contain 8 SECDED bits. Each channel can operate simultaneously, but each must use a different I/O buffer.

Channel Adapters

Channel adapters enable the I/O cluster to communicate with peripheral devices. Several types of channel adapters are available; each type of channel adapter enables the I/O cluster to communicate with a different type of device. Most channel adapters can be connected to only one peripheral device; however, channel adapters for disk storage devices can be connected to multiple devices of the same type.

Each channel adapter corresponds to one I/O buffer. During a data transfer from a peripheral device, the channel adapter converts the input data from the device's format to 64-bit words, generates SECDED bits, and then transmits the data and SECDED bits to the I/O buffer. During a data transfer to an external device, the channel adapter receives 64-bit data words (plus SECDED bits) from the I/O buffer, converts the data to the correct format for that device, and then transmits the data to the device.

Each EIOP controls four channel adapters. The EIOPs control and monitor all data transfers between the I/O buffers and peripheral devices.

The following subsections provide specific information on each type of channel adapter. The specification sheet at the end of this section provides a quick reference summary of all channel adapter specifications.

CCA-1 Channel Adapter

The CCA-1 channel adapter contains a LOSP channel pair that transfers data between an I/O buffer and an external device such as a front-end interface. The LOSP channel pair consists of an input and an output channel. Both channels can operate simultaneously.

Each CCA-1 can support one external device. The maximum transfer rate between the CCA-1 and the external device is 6 or 20 Mbytes/s (software controlled). The word width of each transfer is 16 bits.

All data transfers to or from the CCA-1 are checked for data errors. Data transfers between the CCA-1 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the CCA-1 and peripheral devices are checked for parity errors.

DCA-1 Channel Adapter

The DCA-1 disk channel adapter transfers data between the I/O buffer and a disk drive. The DCA-1 disk channel adapter is compatible with the DD-40, DD-41, and DD-49 disk drives.

Each DCA-1 can support one DD-49 disk drive, two DD-40 disk drives, or two DD-41 disk drives. The maximum transfer rate between the DCA-1 and the disk drive is 12 Mbytes/s. The word width of each transfer is 16 bits.

All data transfers to or from the DCA-1 are checked for data errors. Data transfers between the DCA-1 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the DCA-1 and disk drives are checked for parity errors.

DCA-2 Channel Adapter

The DCA-2 disk channel adapter transfers data between the I/O buffer and a disk drive. The DCA-2 disk channel adapter is compatible with DD-60, DD-61, DD-62, and RD-62 disk drives.

Each DCA-2 can daisy chain up to eight DD-60, DD-61, or DD-62 disk drives; each single daisy chain can support only one drive type. The maximum transfer rate between the DCA-2 and the DD-60 disk drive is 24 Mbytes/s; the maximum transfer rate between the DCA-2 and the DD-61 disk drive is 3 Mbytes/s; and the maximum transfer rate between the DCA-2 and the DD-62 and RD-62 disk drives is 9.34 Mbytes/s. The width of each transfer is 16 bits.

All data transfers to or from the DCA-2 are checked for data errors. Data transfers between the DCA-2 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the DCA-2 and disk drives are checked for parity errors and cyclical redundancy check (CRC) errors.

DCA-3 Channel Adapter

The DCA-3 disk channel adapter transfers data between the I/O buffer and a disk array. The DA-60 comprises five DD-60 spindles; the DA-62 comprises five DD-62 spindles. In each array type, data is striped across four of the spindles, and the fifth spindle is used for odd parity. The DCA-3 disk channel adapter can communicate with up to eight DA-60 or DA-62 disk arrays.

Use of the terms *drive* and *spindle* can be confusing. Whether a device should be referred to as a drive or a spindle is largely determined by the type of channel adapter to which it is connected. If a DD-60 is

connected to a DCA-2 channel adapter, it is an individually accessible I/O device and should be referred to as a drive. However, a DD-60 connected to a DCA-3 represents one-fifth of an array and should be referred to as a spindle.

Disk array performance is basically four times that of the same spindle type connected to a DCA-2.

All data transfers to or from the DCA-3 are checked for data errors. Data transfers between the DCA-3 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the DCA-3 and disk array are checked for parity errors and cyclical redundancy check (CRC) errors.

HCA-3 and HCA-4 Channel Adapters

The HCA-3 and HCA-4 channel adapters enable the IOS-E to communicate with external devices that use a High Performance Parallel Interface (HIPPI) channel. The HIPPI channel pair consists of an input and output channel. Refer to the "High Performance Parallel Interface (HIPPI)" subsection in Section 5 of this manual for more information on the HIPPI channel. The input channel is connected to HCA-3; the output channel is connected to HCA-4.

The HIPPI channel can provide high-speed communications between Cray Research computer systems. The HIPPI channel also enables a Cray Research computer system to be connected to peripheral equipment such as network adapters and graphic display devices.

Each HCA-3 or HCA-4 channel adapter can support one external HIPPI device. The maximum transfer rate between the HCA-3 or HCA-4 and the external device is 100 Mbytes/s. The word width of each transfer is 32 bits.

All data transfers to or from the HCA-3 are checked for data errors. Data transfers from the peripheral equipment to the HCA-3 are checked for parity errors and length/longitudinal redundancy check (LLRC) errors. Data transfers from the HCA-3 to the mainframe or SSD-E are protected by SECDED.

All data transfers to or from the HCA-4 are checked for data errors. Data transfers from the mainframe or the SSD-E to the HCA-4 are protected by SECDED. Data transfers from the HCA-4 to the peripheral equipment are checked for parity errors and LLRC errors.

HCA-5 Channel Adapter

The HCA-5 channel adapter provides an interface between an IOP and an external device. The HCA-5 channel adapter is compatible with any Intelligent Peripheral Interface (IPI) device.

The HCA-5 supports a maximum of one peripheral device. The maximum transfer rate between the HCA-5 and the peripheral device is 213 Mbytes/s. The HCA-5 channel adapter contains eight, bidirectional data buses, which enable its high-transfer rate. The word width of each transfer is 16 bits.

All data transfers to or from the HCA-5 are checked for data errors. Data transfers between the HCA-5 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the HCA-5 and disk drives are checked for parity errors.

TCA-1 Channel Adapter

The TCA-1 tape channel adapter transfers data between the I/O buffer and tape controllers. The TCA-1 channel adapter supports IBM compatible tape controllers.

Each TCA-1 can support up to eight controllers. The maximum number of tape drives each controller supports varies with different tape drive models. The maximum transfer rate between the TCA-1 and a tape drive will also vary with different tape drive models. The word width of each transfer is 8 bits.

All data transfers to or from the TCA-1 are checked for data errors. Data transfers between the TCA-1 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the TCA-1 and peripheral devices are checked for parity errors.

TCA-2 Channel Adapter

The TCA-2 channel adapter provides an interface between an IOP and an external tape controller. The TCA-2 channel adapter is compatible with any Intelligent Peripheral Interface (IPI) device.

The TCA-2 channel adapter is a 16-bit wide interface that uses two bidirectional data busses. Each bus is 8 bits and 1 parity bit (odd parity). The TCA-2 supports a maximum of one peripheral device. The maximum transfer rate between the TCA-2 and the external tape controller is 50 Mbytes/s. The word width of each transfer is 16 bits.

All data transfers to or from the TCA-2 are checked for data errors. Data transfers between the TCA-2 and either the mainframe or the SSD-E are protected by SECDED. Data transfers between the TCA-2 and disk drives are checked for parity errors.

UTC-1 Channel Adapter

The universal time clock channel adapter (UTC-1) quarter board provides resident application programs with read access to the current Greenwich mean time (GMT) and day-of-year (DOY). The UTC-1, accurate to the millisecond, can also notify an application program when a specified GMT has arrived to help control real-time processing. The UTC-1 wiring is implemented only in CRAY C916 series mainframes with serial numbers 4003 and higher.

The UTC-1 receives the time and day information from radio station WWV, a short-wave station operated by the National Institute of Standards and Technology in Boulder, Colorado. WWV transmits the time and date using the universal coordinated time (UCT) standard. The time broadcast by WWV is synchronized with the atomic clocks that provide the time reference for the entire world.

Workstation Interfaces

The two WINs communicate with the workstations: one WIN connects to the operator workstation; the second WIN connects to the maintenance workstation. Each WIN has a 6-Mbyte/s channel pair; each pair consists of an input and an output channel. Each channel contains 16 data bits and 4 parity bits for data error detection.

The workstations communicate with the entire IOS-E through the WINs. Each workstation can send WIN commands that affect the entire IOS-E, a single I/O cluster, or a single I/O processor. The workstations can master clear the entire IOS-E, an individual I/O cluster, or an individual I/O processor. They can also transfer data to or from any IOP, deadstart an IOP, and monitor IOPs for errors.

The IOPs can send requests to the WINs to transfer data to or from a workstation. However, a workstation receiving a request must send a specific command to the requesting IOP before the transfer can begin.

Programmable Real-time Interrupt

The programmable real-time interrupt enables any IOC in the IOS-E to interrupt any CPU in the mainframe. A cable between the mainframe and IOS-E carries the interrupt signals. Each IOC sends 16-bits of parameter data to identify the CPU to be interrupted. Each bit in the parameter data corresponds to a CPU within the mainframe. For example, setting bit 2^1 interrupts CPU 1 in the mainframe which then causes CPU 1 to perform an exchange.

Figure 3-2 is a block diagram that shows the signal paths between the mainframe and IOS-E for the programmable real-time interrupt.

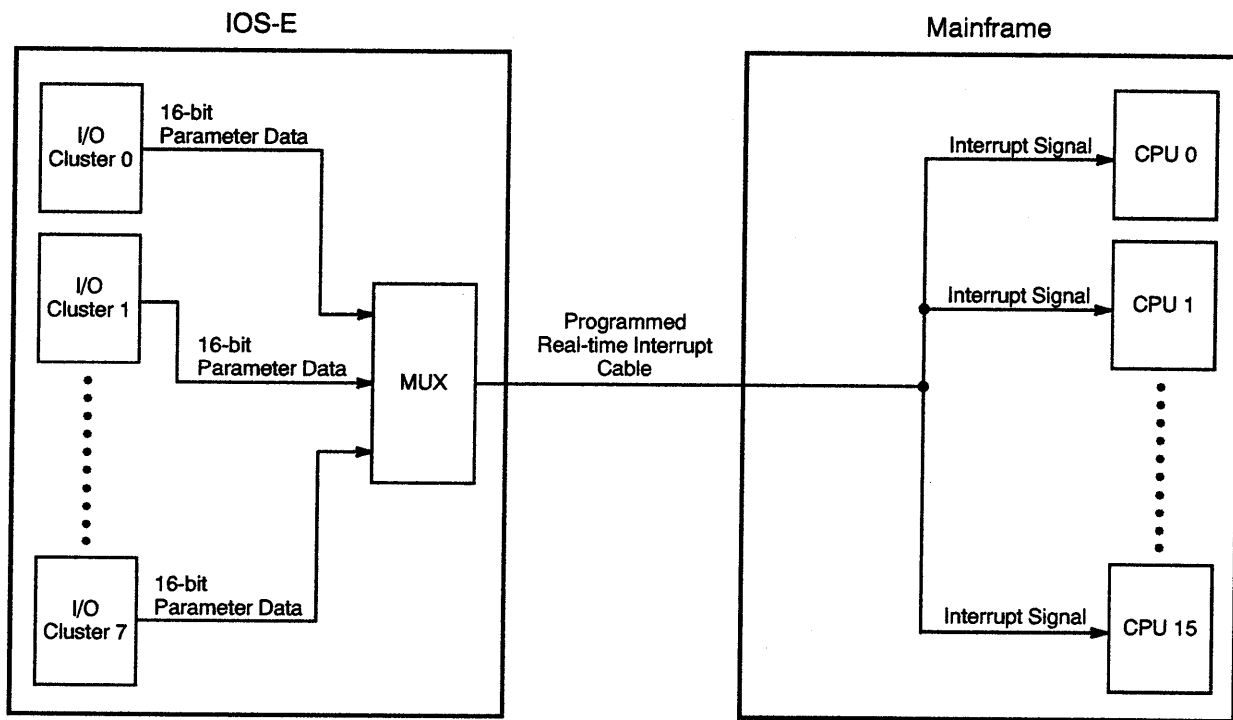



Figure 3-2. Programmable Real-time Interrupt Signal Paths

IOS Model E Specifications



CRAY
RESEARCH, INC.

IOS-E Specifications

General Specifications

I/O clusters 2 to 8
Workstation interfaces 2
Clock speed 160 MHz (6.25 ns)

I/O Cluster

I/O processors:

IOP MUX 1
EIOP 4
Word width 16 bits (1 parcel)
Local memory size 64 Kparcels

I/O buffers 16
Word width 64 bits
Size 64 Kwords
..... (256 Kwords in development)

I/O channels (CPU connection):

Low-speed channel pair 1
Operation full duplex
Channel width 16 bits (1parcel)
Transfer rate 6 Mbytes/s
Data protection parity
High-speed channel pairs 2
Operation full duplex
Channel width 64 bits
Transfer rate 200 Mbytes/s
Data protection SECDED
Channel adapters 16

Channel Adapters

CCA-1:

Operation full duplex
Word width
I/O buffer side 64 bits
External device side 16 bits
Transfer rate 6 or 12 Mbytes/s
Data protection
Mainframe to CCA-1 SECDED
CCA-1 to external interface .. parity
Associated peripheral
devices front-end interfaces
Maximum number of peripheral devices
per CCA-1 1

Channel Adapters (continued)

DCA-1:

Operation half duplex
Word width:
I/O buffer side 64 bits
External device side 16 bits
Transfer rate 12 Mbytes/s
Data protection:
Mainframe to DCA-1 SECDED
DCA-1 to external interface .. parity
Associated peripheral
devices DS-40, DS-41, DS-42
and DD-49 disk drives
Maximum number of peripheral
devices per DCA-1 one DD-49,
one DC-40,
or one DC-41

DCA-2:

Operation half duplex
Word width:
I/O buffer side 64 bits
External device side 16 bits
Transfer rate 24 Mbytes/s
Data protection:
Mainframe to DCA-2 SECDED
DCA-2 to external
interface parity and CRC
Associated peripheral
devices DD-60 and DD-61, or
..... DD-62 disk drives
Maximum number or peripheral
devices per DCA-2 eight DD-60s,
eight DD-61s, eight DD-62s,
or one RD-62

IOS-E Specifications

Channel Adapters (continued)

DCA-3:

Operation half duplex
Word width:
 I/O buffer side 64 bits
 External device side 16 bits
Transfer rate:
 DA-60 80 Mbytes/s
 DA-62 32.56 Mbytes/s
Data protection:
 Mainframe to DCA-3 SECDED
 DCA-3 to external
 interface parity and CRC
Associated peripheral devices ... DA-60,
 or DA-62 disk arrays
Maximum number of peripheral
 devices per DCA-3 eight DA-60s
 or eight DA-62 disk arrays

HCA-3 (HIPPI input channel):

Operation simple duplex
Word width:
 I/O buffer side 64 bits
 External device side 32 bits
Transfer rate 100 Mbytes/s
Data protection:
 Mainframe to HCA-3 SECDED
 HCA-3 to external interface .. parity
Associated peripheral
 devices All HIPPI devices
Maximum number of peripheral
 devices per HCA-3 1

HCA-4 (HIPPI output channel):

Operation simple duplex
Word width:
 I/O buffer side 64 bits
 External device side 32 bits
Transfer rate 100 Mbytes/s
Data protection:
 Mainframe to HCA-4 SECDED
 HCA-4 to external interface .. parity
Associated peripheral
 devices all HIPPI devices
Maximum number of peripheral
 devices per HCA-4 1

HCA-5:

Operation half duplex
Word width:
 I/O buffer side 64 bits
 External device side 16 bits
Transfer rate 213 Mbytes/s
Data protection:
 Mainframe to HCA-5 SECDED
 SSD to HCA-5 SECDED
 HCA-5 to external interface .. parity
Associated peripheral
 devices IPI devices
Maximum number of peripheral
 devices per HCA-5 1

TCA-1:

Operation simple duplex
Word width:
 I/O buffer side 64 bits
 External device side 8 bits
Transfer rate 1.5 to 4.5 Mbytes/s
Data protection:
 Mainframe to TCA-1 SECDED
 SSD to TCA-1 SECDED
 TCA-1 to external device parity
Associated peripheral
 devices IBM compatible tapes
Maximum number of tape
 controllers per TCA-1 8

TCA-2:

Operation half duplex
Word width:
 I/O buffer side 64 bits
 External device side 16 bits
Transfer rate 50 Mbytes/s
Data protection:
 Mainframe to TCA-2 SECDED
 SSD to TCA-2 SECDED
 TCA-2 to external device parity
Associated peripheral
 devices IPI compatible tapes
Maximum number of tape
 controllers per TCA-2 1

IOS-E Specifications

UTC-1:

Operation half duplex

Word width:

External device side 32 bits

Transfer rate:

UTC-1-to-EIOP 50 Mbytes/s

UTC-1-to-mainframe .. 10 Mbytes/s

Data protection:

Mainframe to UTC-1 odd parity

UTC-1 to external

device odd parity

(37 bits of data including odd parity)

Associated peripheral devices:

..... a TRAK microwave interface box,

..... an E-Systems interface box,

..... and a CRAY C916 mainframe

Maximum number of peripherals 3

4 SSD SOLID-STATE STORAGE DEVICES

The Cray Research SSD solid-state storage device model E (SSD-E) and the SSD-E/32i are optional high-performance devices used for temporary data storage. To simplify references to these devices, the term *SSD* used throughout this manual refers to the SSD-E/32i and SSD-E solid-state storage devices, unless specifically stated otherwise.

The SSD functions like a disk drive. Because of its fast access time, fast transfer rates, and large storage capacity, the SSD enhances the performance of a Cray Research computer system by significantly reducing I/O processing time. The storage medium in an SSD is solid-state, dynamic random access memory (DRAM) chips rather than magnetic film. The transfer rates to and from the SSD are considerably faster than that of a disk drive. Data sets for the SSD are identical to data sets for disk drives, providing portability and flexibility.

The mainframe, I/O subsystem model E (IOS-E), and maintenance workstation model E (MWS-E) can transfer data to or receive data from the SSD. The SSD can only respond to transfer requests from these devices; the SSD cannot initiate a transfer.

SSD-E

The following subsection defines the SSD-E physical characteristics, memory, mainframe and SSD-E transfers, IOS-E and SSD-E transfers, and MWS-E and SSD-E transfers. A specification sheet is included at the end of this section.

Physical Description

The SSD-E logic modules can either reside in a separate cabinet from the mainframe or within the mainframe cabinet. The cabinet configuration depends on the memory size. The VHISP channels cannot be configured to support both an external SSD-E and an internal SSD-E. The specification sheet at the end of this section provides the primary physical characteristics of an SSD-E residing in the IOS-E/SSD-E cabinet.

Memory

Each word contains 72 bits: 64 data bits and 8 check bits. The number of words varies, ranging from 128 Mwords to 4 Gwords. Table 4-1 lists the different memory sizes.

Table 4-1. SSD-E Memory Sizes

| Model | Memory Size |
|----------|-------------|
| SSD-128 | 128 Mwords |
| SSD-256 | 256 Mwords |
| SSD-512 | 512 Mwords |
| SSD-1024 | 1 Gword |
| SSD-2048 | 2 Gwords |
| SSD-4096 | 4 Gwords |

To protect data, SSD-E memory uses single-error correction/double-error detection (SECDED) logic.

When data is written into SSD-E memory, the SECDED logic generates a checkbyte (an 8-bit Hamming code †) for each data word. The checkbyte and 64-bit data word are stored in the SSD-E memory.

When a word is read from SSD-E memory, a new checkbyte is generated for the data word. The new checkbyte is compared to the stored checkbyte. The result of the comparison is called the syndrome code. If the syndrome code equals 0, no data bits were altered, and the word is passed on to the external device.

If an error occurred, the SECDED logic analyzes the syndrome code to determine the number of altered data bits. If only a single data bit was altered, the SECDED logic toggles the bit to the correct state and passes the corrected word out to the external device. If two data bits were altered, the SECDED logic cannot correct the word, but it can detect the failure. If more than 2 bits are in error, the results are unpredictable. A message is sent to the error logger for all detected errors.

† Hamming, R. W. "Error Detection and Correcting Codes." *Bell System Technical Journal*. 29.2 (1950): 147 – 160.

Mainframe Data Transfers

Data transfers between the SSD-E and the mainframe's central memory use very high-speed (VHISP) channels. Each VHISP channel simultaneously transfers two 64-bit words plus two 8-bit checkbytes. Each VHISP channel transfers data at the rate of 1,800 Mbytes/s.

The quantity of VHISP channels can range from one to four, depending on the quantity of mainframe CPUs. Typically, there is one VHISP channel for each CPU pair. The CRAY C916 computer system can support eight VHISP channels with 16 CPUs. Because an SSD-E can support 4 VHISP channels, a second SSD-E is required to use all eight VHISP channels from the mainframe.

To protect data, all VHISP channels use SECDED logic. This SECDED logic operates the same as the SECDED logic on the SSD-E memory.

Data transfers between the mainframe and the SSD-E are done in 64-word blocks. Individual words are not accessible by the mainframe. To read a particular word, an entire block is transferred and the word must be selected using software methods similar to disk storage data handling methods.

All VHISP channels operate under mainframe program control. Programming a data transfer requires three parameters: the SSD-E starting address, the mainframe's central memory starting address, and a block length. The block length specifies how many 64-word blocks to transfer. The maximum block length is 16,777,216, which yields a maximum transfer length of 1,073,741,824 words.

IOS-E Data Transfers

Data transfers between the SSD-E and an I/O buffer in the IOS-E use high-speed (HISP) channel pairs. Each HISP channel pair consists of an input and output channel, both of which may be active simultaneously. Each channel is 64 bits wide and contains 8 check bits. Each channel transfers data at the rate of 200 Mbytes/s.

The quantity of HISP channel pairs ranges from two to eight. Typically, there is one HISP channel pair for each I/O cluster in the IOS-E.

To protect data, all HISP channels use SECDED logic. This SECDED logic operates in the same manner as the SECDED logic on the SSD-E memory.

Data transfers between the IOS-E and the SSD-E are done in 64-word blocks. Individual words are not accessible by the IOS-E. To read a particular word, an entire block is transferred and the word must be selected using software methods similar to disk storage data handling methods.

All HISP channel pairs operate under IOS-E program control. Programming a data transfer requires three parameters: the SSD-E starting address, the selected I/O buffer's starting address, and a block length. The block length specifies how many 64-word blocks to transfer. The maximum block length is 1,024, which yields a maximum transfer length of 65,536 words.

MWS-E and SSD-E Transfers

Data transfers between the SSD-E and the MWS-E use the workstation interface (WIN) module located in the IOS and a dedicated 16-bit, 6-Mbyte/s low-speed (LOSP) channel (maintenance interface channel). The maintenance interface channel operates under MWS-E program control and is used for diagnostic maintenance purposes only.

SSD-E/32i

The following subsection defines the SSD-E/32i physical characteristics, memory, mainframe and SSD-E/32i transfers, IOS-E and SSD-E/32i transfers, and MWS-E and SSD-E/32i transfers. A specification sheet is included at the end of this section.

Physical Description

The SSD-E/32i consists of a single-coldplate module located in one of the cabinets containing the computer system. It is installed in a dedicated slot at the top of the mainframe chassis. Standard chassis connections provide power and cooling to the module in a manner similar to all other modules in the chassis. The specification sheet at the end of this section lists the primary physical characteristics of the SSD-E/32i.

Memory

The SSD-E/32i is divided into two independent memory groups: group 0 and group 1. Each group contains a pair of memory banks that can be read from or written to within the same reference. The four block-interleaved memory banks permit simultaneous reading from and/or writing to four different sources. Each memory word contains 72 bits: 64 data bits and 8 check bits. The SSD-E/32i can store 32 Mwords.

To protect data, SSD-E/32i memory uses single-error correction/double-error detection (SECDED) logic. When data is written into SSD-E/32i memory, the SECDED logic generates a checkbyte (an 8-bit Hamming code) for each data word. The checkbyte and 64-bit data word are stored in the SSD-E/32i memory.

When a word is read from SSD-E/32i memory, a new checkbyte is generated for the data word. The new checkbyte is compared to the stored checkbyte. The result of the comparison is called the syndrome code. If the syndrome code equals 0, no data bits were altered, and the word is passed to the external device.

If an error occurs, the SECDED logic analyzes the syndrome code to determine the number of altered data bits. If only a single data bit is altered, the SECDED logic toggles the bit to the correct state and passes the corrected word to the external device. If 2 data bits are altered, the SECDED logic cannot correct the word, but it can detect the failure. If more than 2 bits are in error, the results are unpredictable. A message is sent to the error logger for all detected errors.

Mainframe Data Transfers

Data transfers between the SSD-E/32i and the mainframe's central memory use one very high-speed (VHISP) channel. The VHISP channel operates under mainframe program control. Programming a data transfer requires three parameters: the SSD-E/32i starting address, the mainframe's central memory starting address, and a block length. The block length specifies how many 32-word blocks to transfer.

The SSD-E/32i handles all data in 32-word blocks; all words contain 72 bits. Every data transfer consists of one or more blocks of words. The VHISP channel provides a 144-bit (double-word) path to transfer data to and from SSD-E/32i data buffers. The buffers use 72-bit paths to transfer data to and from SSD-E/32i memory. The maximum number of 32-word blocks that the VHISP channel can transfer in one operation is 1000000₈ (256 Kblocks); the minimum number is 2. The VHISP channel transfer rate is 1,024 Mbytes/s.

IOS-E Data Transfers

Data transfers between the SSD-E/32i and an I/O buffer in the IOS-E use one or two high-speed (HISP) channels pairs. Each HISP channel provides a 72-bit path to and from SSD data buffers. The buffers use 72-bit paths to transfer data to and from SSD memory. The maximum number of 32-word blocks that a HISP channel can transfer in one operation is 100000_8 (32 Kblocks); the minimum number is 1. Each HISP buffer is capable of holding only 1 block and must be loaded or unloaded before another transfer can occur, limiting the HISP channel sustainable transfer rate to 158 Mbytes/s.


To provide data integrity during transfers, both VHISP and HISP channels send a checkbyte (an 8-bit Hamming code) with each 64-bit data word. The checkbyte provides single-error correction/double-error detection (SECDED) during write operations before data is stored. Because the SSD-E/32i stores checkbytes with the data, it is also capable of performing SECDED on data from memory during read operations.

All HISP channel pairs operate under IOS-E program control. Programming a data transfer requires three parameters: the SSD-E/32i starting address, the selected I/O buffer's starting address, and a block length. The block length specifies how many 32-word blocks to transfer.

MWS-E and SSD-E/32i Transfers

Data transfers between the SSD-E/32i and the MWS-E use a workstation interface (WIN) module located in the IOS and a dedicated 16-bit, 6-Mbyte/s low-speed (LOSP) channel (maintenance interface channel). The maintenance interface channel operates under MWS-E program control and is used for diagnostic and maintenance purposes only.

SSD Model E Specifications



SSD-E Specifications

General Specifications

Storage word size 72 bits

Data block size 64 decimal words
per transfer

Maximum block operation:

16,777,215 decimal blocks (7777777₈)
for VHISP channels
256 decimal blocks (377₈)
for HISP channels

Storage capacity:

128 Mwords – 1 Mbit DRAM
256 Mwords – 1 Mbit DRAM
512 Mwords – 1 Mbit DRAM
1 Gword – 4 Mbit DRAM
2 Gword – 4 Mbit DRAM

Memory configurations:

128 Mword - 1 section, 4 groups,
1 Mbit DRAMs
256 Mword - 2 section, 8 groups,
1 Mbit DRAMs
512 Mword - 4 section, 16 groups,
1 Mbit DRAMs
1 Gword - 2 section, 8 groups,
4 Mbit DRAMs
2 Gword - 4 section, 16 groups,
4 Mbit DRAMs
4 Gword - 2 section, 8 groups,
16 Mbit DRAMs

Maximum band width:

100 Gbits/s reading and writing 1 word
per group/clock period

User ports:

4 VHISP channel pairs 1,800 Mbytes/s
each
8 HISP channel pairs 200 Mbytes/s each

Data protection:

Single-error correction/double-error
detection (SECDED) before and after
storage and on all user channels.

SSD-E Physical Description

Height 76.25 in. (194 cm)

Width 46 in. (117 cm)

Depth 73.5 in. (187 cm)

Weight 7,695 lbs (3,182 kg)

Floor loading .. 520 lbs/ft² (2,538 kg/m²)

Access requirements 3 ft (0.9 m)

on all sides

Heat dissipation to air 3.15 kW
(maximum)

SSD-E/32i Specifications

General Specifications

Storage word size 72 bits

Data block size 32 decimal words
per transfer

Maximum block operation:
256 K decimal blocks (1000000₈)
for VHISP channels
32 K decimal blocks (100000₈)
for HISP channels

Storage capacity:
32 Mwords


Memory element:
1 meg × 4 bit DRAM
(1,048,576 4-bit address locations)

Memory bandwidth:
2.214 Gbytes/s, peak

VHISP bandwidth
1 VHISP channel
1,024 Mbytes/s; sustainable, without
contention

HISP bandwidth
2 HISP channel pairs
158 Mbytes/s; sustainable, without
contention

Data protection:
Single-error correction/double-error
detection (SECDED) before and after
storage and on all user channels



5 PERIPHERAL EQUIPMENT

The following subsections describe the major components of the disk drives and various network interfaces used with the CRAY C90 series computer system.

Disk Controller Units and Disk Storage Units

Disk systems provide long-term data storage for a Cray Research computer system. Components of the disk system include disk channel adapters and disk drives. DCA-1 disk channel adapters control one or two DD-40 or DD-41 disk drives. DCA-2 disk channel adapters control from one to eight DD-60, DD-61, or DD-62 disk drives or one RD-62 disk drive. DCA-3 channel adapters control an array of five IPI-2 disk storage units. Refer to Section 3 of this manual for more information on channel adapters.

Disk Drives

The disk drives store data on magnetic disks. Typically, a disk drive consists of several rotating platters. Data is accessed by read and write heads organized into groups. Heads are controlled and positioned by one or more head actuator (servo) mechanisms on the disk cylinders. The following subsections describe specific disk drives.

DD-60 Disk Drive

One DD-60 disk drive consists of a single-sealed head disk assembly (HDA). The HDA contains 2 sets of 9 parallel read and write heads, 1 servo head, 11 eight-inch rotating platters, and 20 thin film media surfaces.

One set of nine parallel heads is used at a time for data transfers. Eight of the heads are used for data and the ninth head is used for parity. The heads can be positioned over 2,608 user-accessible locations on the surface of each platter. Each location is called a cylinder. The DD-60 determines which cylinder the heads are on by reading the information under the servo head.

When the heads are stationary, the area under one head during one complete revolution of the platter is called a track. The DCA-2 disk channel adapter combines eight tracks of data (one from each head) into one logical track. A logical track contains 23 sectors where data can be stored and from which data can be retrieved by the operating system.

Data in one sector is called a data block. One data block consists of 2,048 64-bit words of IOP data plus verification and error-correcting data. Data is transferred between the disk surface and I/O buffer in the IOP in blocks of this fixed size.

One DE-60 disk enclosure cabinet contains a maximum of ten DD-60 disk drives. Eight of the disk drives store system data, and two disk drives are spares. The DCA-2 disk channel adapter in the IOP manages control signals and protocol for the individual disk drives in a DE-60.

The DCA-2 performs the following functions:

- Controls up to eight DD-60 disk drives (daisy chain configuration)
- Passes control functions to the selected drives
- Receives status from the drives
- Generates codes for correcting write data errors
- Checks read data correcting codes and corrects read data when necessary

Initially, a factory flaw table is used to locate media flaws on the surface of a platter. If additional flaws are found, diagnostic programs determine the location and width of the flaw. These flaws, which are identified during surface analysis, are avoided during read and write operations. A defect parameter in the sector ID field contains information on the location of the flaw.

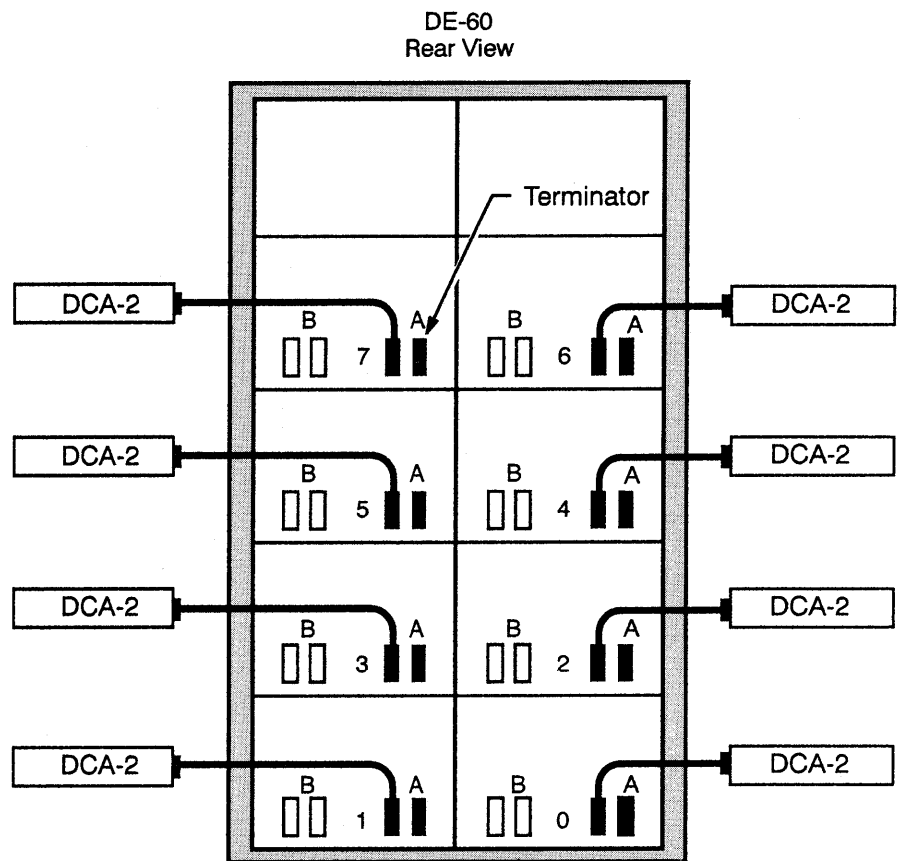
Under control of a DCA-2, a DD-60 writes data into a flawed sector until the media defect location is reached. While the read and write head of the DD-60 is over the media defect, it writes a copy of the previous 18 bytes of data. Then the DD-60 resumes writing valid data in the flawed sector.

When reading data from a flawed sector, a DD-60 reads the defect address to find the beginning of the 18-byte field of repeated data. When the read and write head of the DD-60 reaches the field of repeated data, the DCA-2 does not accept the repeated data. The drive resumes its normal read operation after the head passes the defect field.

DD-60 Single-port Configuration

A single-port configuration connects one DCA-2 to one disk drive. In this configuration, the channel accesses information at the maximum data transfer rate of the disk drive. Because only one disk drive connects to the channel, the storage capacity of the channel is the storage capacity of the disk drive.

Figure 5-1 shows eight disk drives, each connected in a single-port configuration. One DCA-2 connects to the input of port A, and a terminator connects to the output of port A for each disk drive. Port B is not used.



A-10395

Figure 5-1. DD-60 Single-port Configurations

DD-60 Daisy Chain Configuration

A DD-60 daisy chain configuration (refer to Figure 5-2) consists of a maximum of eight DD-60 disk drives connected to one channel. The storage capacity per channel is multiplied by the number of drives connected in the daisy chain; however, only one DD-60 can transfer data to the DCA-2 at a time. Current shipments of 60 series disk drives include a newly designed 2X daisy chain cable. This cable makes it possible for a drive to be removed from a daisy chain without affecting the other units on the chain.

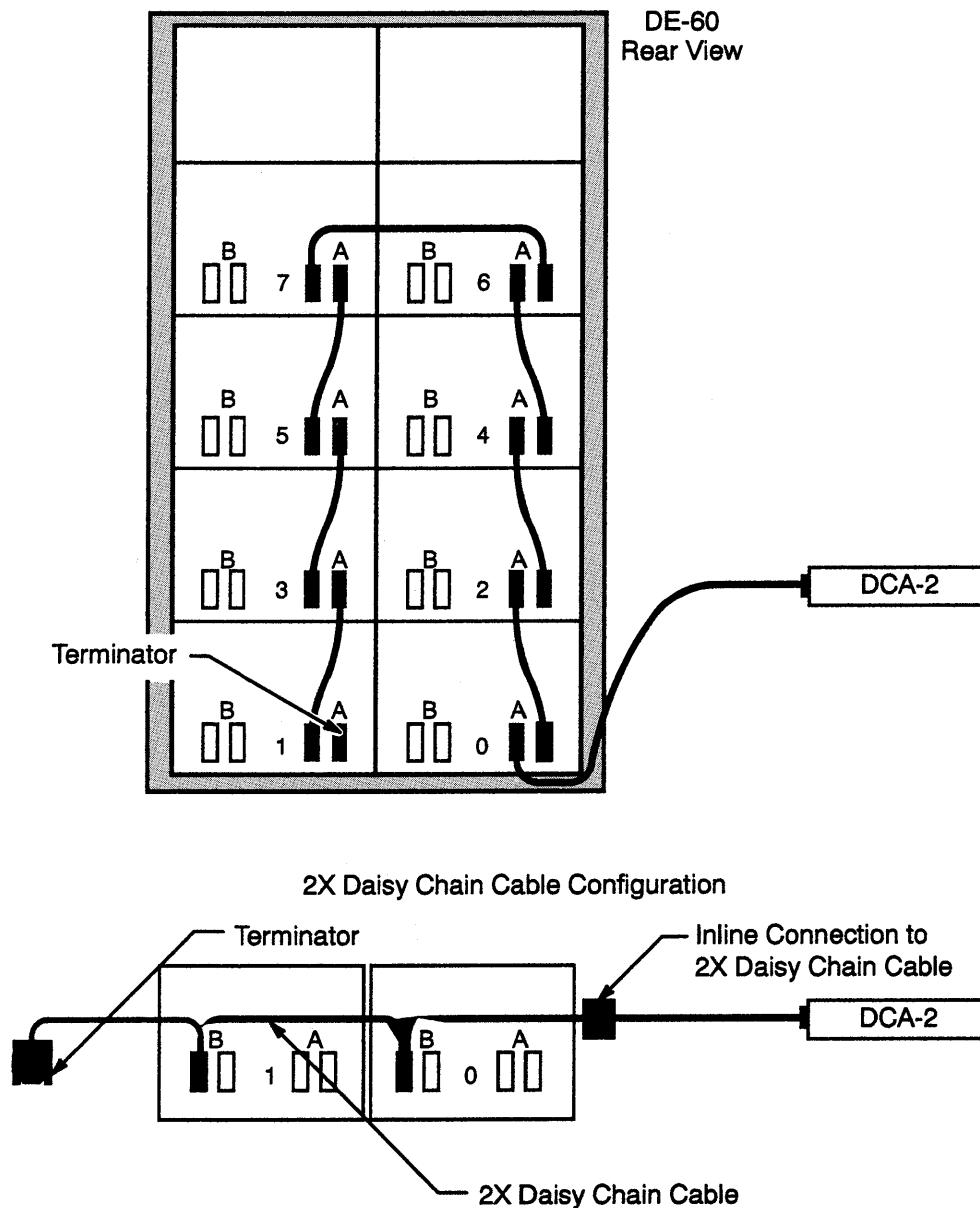


Figure 5-2. DD-60 Daisy Chain Configuration

DD-60 Alternate-path Configuration

An alternate-path configuration connects two DCA-2s to a maximum of eight disk drives. In this configuration, the two DCA-2s connect to the same set of disk drives on two separate daisy chains. Special software modifications must be made when the disk drives are cabled in a redundant configuration.

Figure 5-3 shows eight disk drives connected in redundant configurations. Each disk drive has one DCA-2 connected to port A (primary path) and another DCA-2 connected to port B (redundant path).

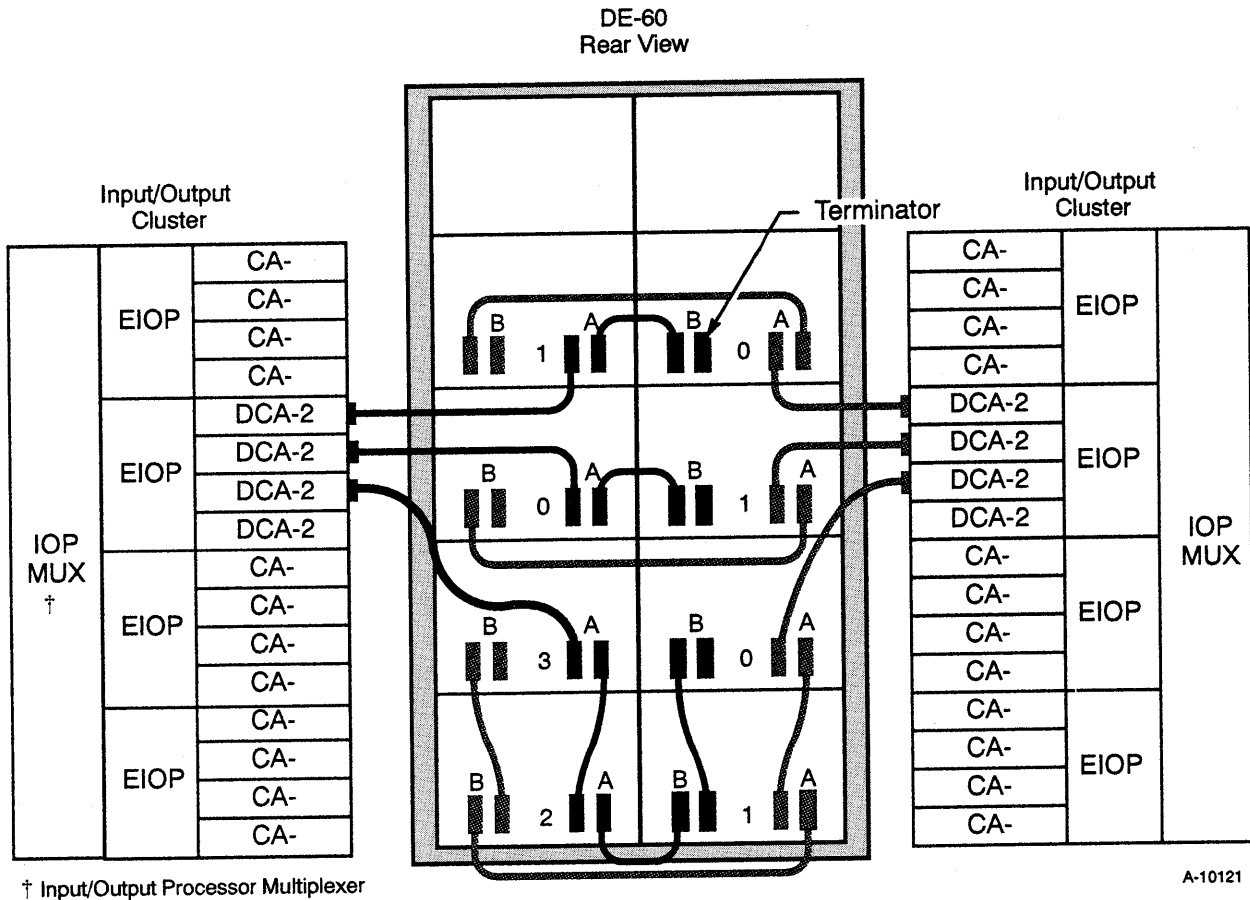


Figure 5-3. DD-60 Alternate-path Configurations

DD-61 Disk Drive

The DD-61 disk drive is a 19-head serial disk drive similar to the DD-60. During data transfers to and from the DCA-2, one head transfers data at a time. The DD-61 has a sustained transfer rate of 2.6 Mbytes/s and a storage capacity of 2.23 Gbytes.

One DD-61 disk drive includes a sealed HDA. The HDA contains 19 serial read and write heads, 1 servo head, 11 eight-inch rotating platters, and 20 thin film media surfaces.

The heads can be positioned over 2,608 user-accessible locations on the surface of each platter. Each location is called a cylinder. The DD-61 determines which cylinder the heads are on by reading the information under the servo head.

When the heads are stationary, the area under one head after one complete revolution of the platter is called a track. A track contains 11 sectors where data can be stored and from which data can be retrieved by the operating system.

Data in one sector is called a data block. One data block consists of 512 64-bit words of IOP data plus verification and error-correcting data. Data is transferred between the disk surface and the I/O buffer in the IOP in blocks of this fixed size. Sectors may be chained during both read and write operations.

One DE-60 disk enclosure cabinet contains a maximum of ten DD-61 disk drives. Eight of the disk drives store system data, and two disk drives are spares. The DCA-2 disk channel adapter in the IOP manages control signals and protocol for the individual disk drives in a DE-60.

The DCA-2 performs the following functions:

- Controls a maximum of eight DD-61 disk drives (daisy chain configuration)
- Passes control functions to the selected drives
- Receives status from the drives
- Generates codes for correcting write data errors
- Checks read data correcting codes and corrects read data when necessary

Initially, a factory flaw table is used to locate media flaws on the surface of a platter. If additional flaws are found, diagnostic programs determine the location and width of the flaw. These flaws, which are identified during surface analysis, are avoided during read and write operations. A defect parameter in the sector ID field contains information on the location of the flaw.

Under control of a DCA-2, a DD-61 writes data into a flawed sector until the media defect location is reached. While the read and write head of the DD-61 is over the media defect, it writes a copy of the previous 18 bytes of data. Then the DD-61 resumes writing valid data in the flawed sector.

When reading data from a flawed sector, a DD-61 reads the defect address to find the beginning of the 18-byte field of repeated data. When the read and write head of the DD-61 reaches the field of repeated data, the DCA-2 does not accept the repeated data. The drive resumes its normal read operation after the head passes the defect field.

DD-61 disk drives also connect in the same single-port, daisy chain, or alternate-path configurations as DD-60s. For information on the daisy chain and alternate-path configurations, refer to the "DD-60 Disk Drive" subsection in this section.

DD-62 Disk Drive

The DD-62 is a two-head parallel storage unit. It has a sustained transfer rate of 8.14 Mbytes/s and a storage capacity of 2.73 Gbytes. One DD-62 contains nine read/write head groups and one servo head. During data transfers to and from the DCA-2, two heads are used at a time. The servo head transfers head-position information to the servo control circuitry in the DD-62.

A sector of data from a DD-62 contains 512 64-bit words of IOP data. Data is transferred between the DD-62 and DCA-2 in blocks of this fixed size. Each track contains 28 sectors where data can be stored and retrieved by the IOS-E.

The DCA-2 creates a sector from two physical sectors in the DD-62 (one from each head in the head group). Each physical sector contains one half of an IOP data sector. Nine logical tracks make up one cylinder in the DD-62. DD-62s contain 2,652 data cylinders, 2 maintenance cylinders, and 1 flaw table cylinder.

One DE-60 disk enclosure cabinet contains a maximum of ten DD-62 disk drives. Eight of the disk drives store system data, and two disk drives are spares.

RD-62 Disk Drive

The RD-62 is a two-head parallel storage unit that is identical to the DD-62 in performance. It has a sustained transfer rate of 8.14 Mbytes/s and a storage capacity of 2.73 Gbytes.

The RD-62 is housed in an RDE-6 enclosure that enables individual drives to be easily removed and replaced by the customer. The RDE-6 enclosure contains up to four RD-62s. Connections to the RD-62s are made through a bulkhead on the RDE-6 cabinet. Because of the limitations of the RDE-6 bulkhead, RD-62s do not support daisy chain or alternate-path cabling configurations like the DD-62s. In all other respects, the RD-62 is equivalent to the DD-62.

Disk Array

The DCA-3 channel adapter supports a five-spindle disk array composed of DD-60 or DD-62 spindles. Four of the spindles hold data, and the fifth spindle contains parity information on the data. The spindles are housed in DE-60 disk enclosure cabinets; each cabinet contains up to ten spindles. As many as eight disk arrays can be daisy chained on a single DCA-3. Figure 5-4 is an overview of a two-array daisy chain.

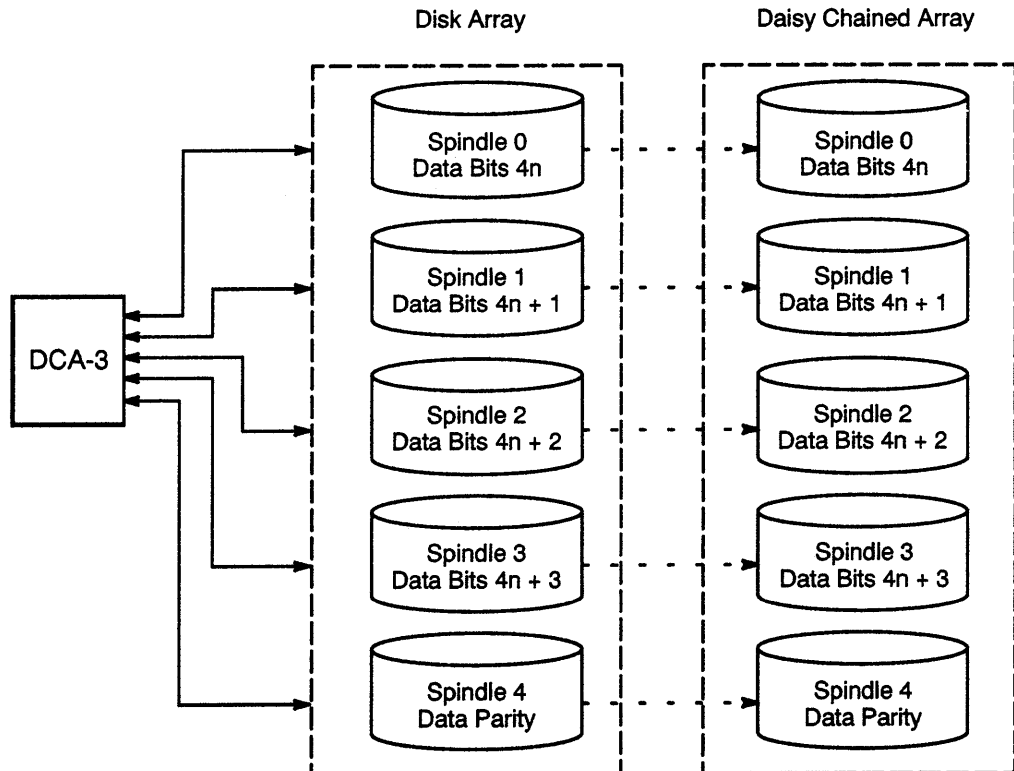


Figure 5-4. Disk Array Overview Block Diagram

DS-41 Disk Subsystem

The DS-41 disk subsystem consists of the DC-41 disk controller and the DD-41 disk drive. Each DD-41 disk drive has four spindles that operate as a single logical disk drive unit under the control of one DC-41 disk controller. Each physical disk drive (spindle) consists of 9 rotating platters and 15 recording surfaces. Data is accessed by 15 read and write heads. A servo mechanism, which controls the read and write heads, positions the heads over one of 1,635 disk cylinders. Data is stored and retrieved from the recording surface of the disk drive by any of the 15 read and write heads.

The recording surface available to each head is called a disk track, which is the basic storage unit reserved by the operating system. Each disk track has 48 sectors where data can be stored and retrieved by the operating system. The data in one sector is called a data block. One data block consists of 2,048 16-bit parcels (512 64-bit words) of IOP data plus verification and error-correction data. Data can be transferred between the disk surface and local memory in the IOP only in blocks of this fixed size. Sectors may be chained for both read and write operations.

A DC-41 disk controller provides interface logic to adapt DCA-1 signals and protocol for individual disk drive units, to handle routing among the drives, and to buffer data from the four spindles in a full-track buffer. The interface logic in one DC-41 disk control unit performs the following functions:

- Controls up to 8 spindles (two DD-41 disk drives)
- Passes control functions to the selected drives
- Passes status from the drives to the DCA-1
- Buffers read and write data for transfers between DCA-1 and the disk drives
- Generates error-correction codes for write data
- Checks read data correction codes and corrects read data when necessary
- Controls distribution of read and write data over 48 sectors per track using 12 sectors from each of the four spindles in a logical drive

Under the control of a DC-41, a disk drive writes data to a flawed sector until a defect location is reached. In the area starting at a defect address, a disk drive writes a 16-byte field of 0's. A disk drive resumes writing data in a flawed sector following this field of 0's. When reading data from a flawed sector, a disk drive reads the defect address to find where the field of 16 bytes of 0's starts. When a drive's read and write head reaches the field of 0's, the head skips over the flawed area of the sector overwritten by the field of 0's, omitting them from the read data. The drive resumes its normal read operation after the head passes the defect field.

A factory flaw table is used initially; if any additional flaws are found, diagnostic programs determine where the flaw is located in the sector and how wide it is. Defective areas of the recording surface, which are identified during surface analysis, are avoided during read and write operations. A defect parameter becomes part of the sector ID field when the drive is formatting.

DS-41 Single-port Configuration

A single-port configuration connects one DC-41 to one DD-41. In this configuration, the channel accesses information at the maximum data-transfer rate of the DD-41. Because only one disk drive connects to the channel, the storage capacity of the channel is the storage capacity of the disk drive.

DS-41 Daisy Chain Configuration

A daisy chain configuration connects one DC-41 to two DD-41s. The channel data-storage capacity is the total storage capacity of both disk drives. Because only one disk drive can transfer data to the DC-41 at a time, the channel data transfer rate is the maximum transfer rate of one disk drive.

DS-41 Alternate-path Configuration

An alternate-path configuration connects two DC-41s to a maximum of two disk drives. In this configuration, the two DC-41s connect to the same disk drives on two separate daisy chains. Special software modifications must be made when the disk drives are cabled in an alternate-path configuration.

DS-41A Disk Subsystem Field-upgradable Configurations

A field-upgradable DS-41A disk subsystem configuration has the following components:

- One disk drive cabinet (DE-41)
- One DD-41 disk drive (housed in the DE-41)
- One disk controller cabinet (DCC-2A)
- One DC-40 disk controller (housed in the DCC-2A)

A DS-41A can be upgraded by adding a DS-41B package. A DS-41B consists of one DD-41, one DC-41, and all cabling required to install the additional drive and controller in a DS-41A disk subsystem. Up to three DS-41Bs can be installed in a DS-41A disk subsystem.

DS-40 Disk Subsystem

The DS-40 disk subsystem comprises the following components: the DD-40 disk drive, the DC-40 disk control unit (DCU), and the disk controller cabinet (DCC-2). The DD-40 contains four disk drives and required interface logic to operate as a single disk drive unit. The DC-40 is housed in the DCC-2, which is separate from the DD-40 disk drives. Refer to "DS-40 and DS-40D Disk Subsystem Specifications" at the end of this section for exact configuration information. Each physical disk drive (spindle) consists of six rotating platters and ten recording surfaces. Data is accessed by 19 read and write heads that are controlled and positioned by a servo mechanism to one of 1,418 disk cylinders.

The recording surface available to each head is called a disk track, which is the basic storage unit reserved by the operating system. Each disk track has 48 sectors where data can be recorded and read. The data in one sector is called a data block. One data block consists of 2,048 16-bit data parcels (512 64-bit words) plus verification and error-correction data. Data can be transferred between the disk surface and I/O buffer in the IOP only in blocks of this fixed size. Sectors may be chained for both read and write operations.

Interface logic in the DC-40 also adapts the DCA-1 signals and protocol to the individual disk drive units, manages routing among the drives, and buffers the data from the four drives in a full-track buffer.

The interface logic in one DC-40 disk control unit performs the following functions:

- Controls up to 8 spindles (two DD-40 disk storage units)
- Passes control functions to the selected drives
- Passes status from the drives to the DCA-1
- Buffers read and write data for transfers between DCA-1 and the disk drives
- Generates error-correction codes for write data
- Checks read data correction codes and corrects read data when necessary
- Controls distribution of read and write data over 48 sectors per track using 12 sectors from each of the four spindles in a logical drive

Under the control of a DC-40, a disk drive writes data onto a flawed sector until a defect location is reached. In the area starting at a defect address, a disk drive writes a 16-byte field of 0's. A disk drive resumes writing data in a flawed sector following this field of 0's. When reading data from a flawed sector, a disk drive reads the defect address to find where the field of 16 bytes of 0's starts. When a drive's read and write head reaches the field of 0's, the head ignores the field of 0's, omitting the field of 0's from the read data. The drive resumes its normal read operation after the head passes the defect field.

A factory flaw table is used initially; if any additional flaws are found, diagnostic programs determine where the flaw is located in the sector and how wide it is. Defective areas of the recording surface, which are identified during surface analysis, are avoided during read and write operations. A defect parameter becomes part of the sector ID field when the drive is formatting.

DS-40 Single-port Configuration

A single-port configuration connects one DC-40 to one DD-40. In this configuration, the channel accesses information at the maximum data-transfer rate of the DD-40. Because only one disk drive connects to the channel, the storage capacity of the channel is the storage capacity of the disk drive.

DS-40 Daisy Chain Configurations

A daisy chain configuration connects one DC-40 to two DD-40s. The channel data-storage capacity is the total storage capacity of both disk drives. Because only one disk drive can transfer data to the DC-40 at a time, the channel data-transfer rate is the maximum transfer rate of one disk drive.

DD-49 Disk Drive

The DD-49 disk drive consists of nine rotating platters. Data is accessed by 32 read and write heads organized into eight groups, four read and write heads per group. Heads are controlled and positioned by two identical head actuator (servo) mechanisms to one of 886 disk cylinders. The servo mechanisms are identified as Servo-A and Servo-B.

The recording surface available to each head group is called a disk track, and is the basic storage unit reserved by the operating system. Each disk track has 42 sectors (and two spare sectors) where data is recorded and read back. The data in one sector is called a data block and consists of 2,048 16-bit parcels (512 64-bit words) of IOP data plus verification and error-correction data. Data can be transferred between the disk surface and I/O buffer in the IOP only in blocks of this fixed size. Sectors may be chained for both read and write operations.

The DD-49 disk drive responds to commands from the IOS-E through a microprocessor unit card (MPU card) that contains a 68000-type 16-bit microprocessor and a second processor called the supervisor.

The DD-49 disk drive provides a sector-slipping mechanism that allows a full track to remain available to the system even after one or two sectors of the track become flawed. Sectors are slipped from the flawed sector to the end of the track. In general, if sector *n* becomes flawed, sectors *n* through 41 of the track are slipped, and the data contained in sectors *n* through 41 must be re-created. If a second sector in a track becomes flawed, the process is repeated. If a third sector in a track becomes flawed, the operating system must mark the sector as unavailable. Sector slipping takes place offline. A hardware diagnostic utility reformats the track that contains slipped sectors.

A DD-49 disk drive has 44 sectors per track, although only 42 sectors are used for data. Under normal circumstances, the two spare sectors are not used. If one of the data sectors becomes flawed, however, a spare sector is used as a data sector.

Refer to "DD-49 Disk Drive Specifications" at the end of this section for configuration information.

Network Interfaces

The CRAY C90 series computer system can be connected to a wide variety of computer systems (often referred to as “front-end systems”) and networks through a CCA-1 channel adapter in the IOS-E. This enables users of non-Cray Research computer systems to use the extraordinary computational power of the CRAY C90 series system. The following subsections describe the methods used to interface the CRAY C90 series computer system with other computer systems and networks.

FEI-1 Front-end Interface

The FEI-1 front-end interface provides communication between a CCA-1 channel adapter in the IOS and many different types of front-end computer systems. The FEI-1 compensates for differences in channel widths, machine word size, electrical logic levels, and control protocols. Refer to “Front-end Interface Specifications” at the end of this section for a complete list of compatible mainframes and minicomputers.

The FEI-1 is housed in a stand-alone cabinet located near the host computer. The cabinet is air cooled and operates directly from the AC power mains; power consumption varies with each type of interface. Internal power supplies provide all required voltages. Cabinet grounding is flexible and can be configured to specific site requirements.

Each FEI-1 contains two or more logic modules and the appropriate cabling. The hardware logic contained in these modules performs all command translation and protocol conversion needed to transfer data; these operations are invisible to both the front-end and Cray Research programmer.

Fiber-optic Link

The Cray Research fiber-optic link (FOL) is used as a channel extender for 6-Mbyte/s (16-bit asynchronous) channels. It connects the conventional wire cable from the CCA-1 channel adapter in the IOS to the wire cable from an FEI-1. Fiber-optic cabling enhances the performance of the FEI-1 by eliminating the occasional problems related to system isolation, including induced noise, variable ground potentials, and radio frequency radiation found in wire cabling. Fiber-optic cabling overcomes these problems and, in addition, provides a secure link for transmitting data over distances up to 4,000 m.

Fiber-optic technology uses thin glass fibers (optical fibers) to transmit information from one location to another. Optical fibers are used in place of wire cabling, and light signals replace electrical charges sent over conventional wire cabling.

The FOL operates by converting digital data into electrical pulses. The electrical signal is used to modulate light coming from a light-emitting diode (LED). The resulting light pulses, which are of the same duration as electrical pulses, are sent over the fiber-optic cable. At the receiving end, the light pulses are converted back into electrical pulses, which are then demodulated to recover the digital data. As with a standard FEI-1, these operations are invisible to both the front-end and Cray Research programmer.

The fiber-optic FEI-1 cabinet is similar to the standard FEI-1 cabinet. It is modified with an attached compartment to hold the fiber-optic modules. In addition to this FEI-1 cabinet, another smaller cabinet containing more fiber-optic modules is located next to the Cray Research mainframe. These special fiber-optic modules modulate and demodulate the signals between the Cray Research mainframe and the front-end system.

FEI-3 Front-end Interface

The FEI-3 is a group of similar front-end interfaces that enables VME-based microcomputers and workstations to communicate with a CCA-1 channel adapter in the IOS over a standard 6-Mbyte/s I/O channel. Specific FEI-3 applications depend on the capabilities of the VME workstations or microcomputers. For example, Cray Research uses the FEI-3 to connect systems to an operator workstation.

The following list contains other possible FEI-3 applications:

- To connect to a communications gateway for Control Subsystem Networks or other networks
- To connect to a graphics output processor or device
- To connect to a remote Cray Research station

Each FEI-3 interface consists of two VME-compatible circuit boards that install into the target VME system, plus supporting cables and software drivers. The customer furnishes and provides support for the target VME system.

The VMEbus is an industry standard that specifies the electrical and mechanical rules for a microcomputer backplane. Many popular microcomputer systems are based on the VMEbus.

Direct Network Connections

The CCA-1 channel adapter in the IOS supports direct connection to network adapters such as Network Systems Corporation HYPERchannel adapters, Computer Network Technology Corporation LANlord adapters, and others.

High Performance Parallel Interface (HIPPI)

The Cray Research High Performance Parallel Interface (HIPPI) is an external channel that provides high-speed communications between HCA-3 and HCA-4 channel adapters in the IOS and peripheral equipment, such as network adapters, raster display devices, and mass storage systems. HIPPI conforms to industry standards and provides 32-bit parallel data transfers at the rate of 100 Mbytes/s.

HIPPI conforms to the preliminary draft proposed American National Standard (DPANS) HIPPI revision 7.0. The HIPPI proposal is based on an original design by engineers at Los Alamos National Laboratories.

HIPPI is a simplex channel that transmits data in one direction; it is usually configured in pairs for full duplex operation. Driver software enables users to operate the HIPPI directly as a raw device or indirectly through Transmission Control Protocol/Internet Protocol (TCP/IP) or User Datagram Protocol (UDP) sockets, Remote Procedure Call (RPC) libraries, and Network File Systems (NFSs) between Cray Research computer systems.

Because HIPPI conforms to industry standards, it can be configured with many types of devices and applications that require high-speed transfer of large amounts of data.


The following list contains other HIPPI applications:

- Distributed applications. The speed of HIPPI makes more applications suitable for distributed processing. Users can link multiple Cray Research computer systems for maximum supercomputer performance.
- Raster graphics. Real-time animated graphics are possible when HIPPI is combined with a compatible high-speed frame buffer. Existing devices have delivered up to 60 frames per second on a 512-by-512 raster of 24-bit pixels.

DEC VAX Supercomputer Gateway

Digital Equipment Corporation (DEC) offers a VAX Supercomputer Gateway to enable direct connection between the DEC VAXcluster environment and a CCA-1 channel adapter in the IOS.

DD-60 Disk Drive Specifications



DD-60 Specifications

DD-60 Features

Transfer rate:
Sustained 16 to 20 Mbytes/s
Burst rate 24 Mbytes/s

Storage capacity
One DD-60 1.96 Gbytes

Total data sectors 119,968

Logical sector size (64-bit words) 2,048

Total data words 245,694,464

Typical position delays:
Single track 3 ms
Average 13 ms
Full stroke 26 ms

Average latency 8.3 ms

DE-60 Power and Cooling Specifications (ten DD-60s)

Required power 200 to 208 Vac, 3 phase,
50 or 60 Hz, 12 A per phase
or
380 to 416 Vac, 3 phase,
50 or 60 Hz, 7 A per phase

Heat load
(8 disk drives) .. 8,600 Btu/hr, (2,520 W)

Type of cooling air cooled

DE-60 Physical Description

Dimensions
Height 61.7 in. (157 cm)
Width 24.0 in. (61 cm)
Depth 41.5 in. (105 cm)

Floor space 6.9 ft² (0.6 m²)

Weight
(all ten DD-60s) 960 lbs (435 kg)


DE-60 Placement and Cabling Specifications

Minimum clearance
Sides 2 in. (5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of
data cables 98.4 ft (30 m)

DD-61 Disk Drive Specifications



DD-61 Specifications

DD-61 Features

Transfer rate:

Sustained 2.3 to 2.6 Mbytes/s
Burst rate 3.0 Mbytes/s

Storage capacity

One DD-61 2.23 Gbytes

Total data sectors 545,072

Logical sector size (64-bit words) 512

Total data words 279,076,864

Typical position delays:

Single track 3 ms
Average 13 ms
Full stroke 26 ms

Average latency 8.3 ms

DE-60 Power and Cooling Specifications (ten DD-61s)

Required power 200 to 208 Vac, 3 phase,
50 or 60 Hz, 6.4 A per phase
or
380 to 416 Vac, 3 phase,
50 or 60 Hz, 3.7 A per phase

Heat load

(8 disk drives) ... 4,770 Btu/hr, (1,400 W)

Type of cooling air cooled

DE-60 Physical Description

Dimensions

Height 61.7 in. (157 cm)
Width 24.0 in. (61 cm)
Depth 41.5 in. (105 cm)

Floor space 6.9 ft² (0.6 m²)

Weight

(all ten DD-61s) 812 lbs (368 kg)

DE-60 Placement and Cabling Specifications

Minimum clearance


Sides 2 in. (5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of

data cables 98.4 ft (30 m)

DD-62 Disk Drive Specifications



DD-62 Specifications

DD-62 Features

Transfer rate:
Sustained 8.14 Mbytes/s
Burst rate 9.34 Mbytes/s

Storage capacity
One DD-62 2.73 Gbytes

Total data sectors 545,072

Logical sector size (64-bit words) 512

Total data words 279,076,864

Typical position delays:
Single track 3 ms
Average 12 ms
Full stroke 26 ms

Average latency 6.87 ms

DE-60 Power and Cooling Specifications (ten DD-62s)

Required power 200 to 208 Vac, 3 phase,
50 or 60 Hz, 6 A per phase
or
380 to 416 Vac, 3 phase,
50 or 60 Hz, 3 A per phase

Heat load
(8 disk drives) ... 5,700 Btu/hr, (1,670 W)

Type of cooling air cooled

DE-60 Physical Description

Dimensions
Height 61.7 in. (157 cm)
Width 24.0 in. (61 cm)
Depth 41.5 in. (105 cm)

Floor space 6.9 ft² (0.6 m²)

Weight
(all ten DD-62s) 810 lbs (367 kg)


DE-60 Placement and Cabling Specifications

Minimum clearance
Sides 2 in. (5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of
data cables 98.4 ft (30 m)

RD-62 Disk Drive Specifications



CRAY
RESEARCH, INC.

RD-62 Specifications

RD-62 Features

Transfer rate:
Sustained 8.14 Mbytes/s
Burst rate 9.34 Mbytes/s

Storage capacity
One RD-62 2.73 Gbytes

Total data sectors 545,072

Logical sector size (64-bit words) 512

Total data words 279,076,864

Typical position delays:
Single track 3 ms
Average 12 ms
Full stroke 26 ms

Average latency 6.87 ms

RDE-6 Power and Cooling Specifications (four RD-62s)

Required power 208 to 240 Vac, 1 phase,
50 or 60 Hz, 6 A

Heat load 2,460 Btu/hr, (720 W)

Type of cooling air cooled

RDE-6 Physical Description

Dimensions
Height 42.0 in. (107 cm)
Width 23.0 in. (58 cm)
Depth 36.0 in. (91 cm)

Floor space 5.8 ft² (0.5 m²)

Weight
(all four RD-62s) 494 lbs (224 kg)


RDE-6 Placement and Cabling Specifications

Minimum clearance
Sides 1 in. (2.5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of
data cables 98.4 ft (30 m)

DA-60 Disk Drive Specifications



DA-60 Specifications

DA-60 Features

Transfer rate:
Sustained 64 to 80 Mbytes/s
Burst rate 96 Mbytes/s

Storage capacity
One DA-60 7.84 Gbytes

Total data sectors 119,968

Logical sector size (64-bit words) 8,192

Total data words 982,777,856

Typical position delays:
Single track 3 ms
Average 13 ms
Full stroke 26 ms

Average latency 8.3 ms

DA-60 Power and Cooling Specifications (ten DD-60s)

Required power 200 to 208 Vac, 3 phase,
50 or 60 Hz, 6 A per phase
or
380 to 416 Vac, 3 phase,
50 or 60 Hz, 3 A per phase

Heat load
(8 disk drives) .. 8,600 Btu/hr, (2,520 W)

Type of cooling air cooled

DA-60 Physical Description

Dimensions
Height 61.7 in. (157 cm)
Width 24.0 in. (61 cm)
Depth 41.5 in. (105 cm)

Floor space 6.9 ft² (0.6 m²)

Weight
(all ten DD-60s) 960 lbs (435 kg)


DA-60 Placement and Cabling Specifications

Minimum clearance
Sides 2 in. (5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of
data cables 98.4 ft (30 m)

DA-62 Disk Drive Specifications



DA-62 Specifications

DA-62 Features

Transfer rate:
Sustained 32.5 Mbytes/s
Burst rate 37.36 Mbytes/s

Storage capacity
One DA-62 10.92 Gbytes

Total data sectors 545,072

Logical sector size (64-bit words) 2,048

Total data words 1,116,307,456

Typical position delays:
Single track 3 ms
Average 12 ms
Full stroke 26 ms

Average latency 6.87 ms

DA-62 Power and Cooling Specifications (five DA-62s)

Required power 200 to 208 Vac, 3 phase,
50 or 60 Hz, 6 A per phase
or
380 to 416 Vac, 3 phase,
50 or 60 Hz, 3 A per phase

Heat load
(8 disk drives) ... 5,700 Btu/hr, (1,670 W)

Type of cooling air cooled

DA-62 Physical Description

Dimensions
Height 61.7 in. (cm)
Width 24.0 in. (cm)
Depth 41.5 in. (cm)

Floor space 6.9 ft² (0.6 m²)

Weight
(all ten DD-62s) 810 lbs (367 kg)



DA-62 Placement and Cabling Specifications

Minimum clearance
Sides 2 in. (5 cm)
Front 36 in. (91 cm)
Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of
data cables 98.4 ft (30 m)

DS-40 and DS-40D Disk Subsystem Specifications



CRAY
RESEARCH, INC.

DS-40 and DS-40D Specifications

DC-40 Features

Transfer rate:
Sustained 9.6 Mbytes/s
Burst rate 20 Mbytes/s
Storage capacity 5,200 Mbytes

DC-40 Power and Cooling

Required power 208 Vac, 3 phase,
60 Hz, 60 A
Type of cooling water cooled
refrigeration/air cooling
Water temperature (°F) 40 to 90
Water temperature (°C) 4.4 to 32.2
Heat load (to air) 1,330 Btu/hr, 390 W
Heat rejection
to water 24,000 Btu/hr, 7,643 W

DCC-2/DC-40

Physical Description

The four DC-40s are housed in a disk control cabinet (DCC-2) that contains the power control and refrigeration components required for the DC-40.

Floor space 8.7 ft² (0.81 m²)
Weight 1,240 lbs (562 kg)
Cabinet dimensions:
Height 60 in. (152 cm)
Width 31 in. (79 cm)
Depth 41 in. (104 cm)

DCC-2/DC-40

Placement and Cabling

Minimum clearance:
Sides 12 in. (30.5 cm)
Front 36 in. (91.4 cm)
Back 36 in. (91.4 cm)
Length of power cable 8 ft (2.4 m)
Maximum length of
data cables 50 ft (14.4 m)

DD-40 Features

Transfer rate:
Sustained 9.6 Mbytes/s
Burst rate 20 Mbytes/s

DD-40 Features (continued)

Total data sectors 1,293,216
Total data words 662,126,592
Typical position delays:
Single track 4 ms
Average 16 ms
Full stroke 30 ms

DD-40 Power and Cooling

Required power 208 Vac, 3 phase,
60 Hz, 20 A
Cooling air cooled
Heat load 8,000 Btu/hr, 2,340 W

DD-40 Physical Description

Floor space 7.3 ft² (0.68 m²)
Weight 1,150 lbs (522 kg)
Cabinet dimensions:
Height 60 in. (152 cm)
Width 26 in. (66 cm)
Depth 41 in. (104 cm)

DD-40 Placement and Cabling

Minimum clearance:
Sides 1 in. (2.5 cm)
Front 36 in. (91.4 cm)
Back 30 in. (76.2 cm)



Length of power cable 6 ft (1.8 m)
Maximum length of data cables ... 20 ft (6 m)

The DCC-2 contains four DC-40 disk controllers. The DC-40 is a dual-ported interface with only one port active at a time.

Four disk storage units (DSUs) are connected to the DCC-2 chassis for the DS-40 Disk Subsystem.

Eight DSUs are connected to the DCC-2 chassis for the DS-40D disk subsystem, but only four DSUs can be active at one time. This technique, known as daisy chaining, is used to double the capacity of a single subsystem from 21 Gbytes to 42 Gbytes; doubling the capacity does not double the performance because the data path is set at 9.6 Mbytes/s.

DS-41, DS-41D, and DS-41R Disk Subsystem Specifications



DS-41, DS-41D, and DS-41R Specifications

DC-41 Features

Sustained transfer rate 9.6 Mbytes/s
Storage capacity 4,800 Mbytes

DCC-2A Power and Cooling

(four DC-41s)

Required power 208 Vac, 3 phase,
60 Hz, 60 A

Type of cooling water cooled
refrigeration/air cooling

Water temperature (°F) 40 to 90

Water temperature (°C) 4.4 to 32.2

Heat load (to air) 1,330 Btu/hr, 390 W

Heat rejection
to water 24,000 Btu/hr, 7,643 W

DCC-2A Physical Description

(four DC-41s)

The DC-41s are housed in a disk control cabinet (DCC-2A) that contains the power control and refrigeration components required for the DC-41.

Floor space 8.7 ft² (0.81 m²)

Weight 1,240 lbs (562 kg)

Cabinet dimensions:

Height 67 in. (170 cm)

Width 31 in. (79 cm)

Depth 41 in. (104 cm)

DCC-2A Placement and Cabling

Minimum clearance:

Sides 12 in. (30.5 cm)

Front 36 in. (91.4 cm)

Back 36 in. (91.4 cm)

Length of power cable 8 ft (2.4 m)

Maximum length of
data cables 50 ft (14.4 m)

The DCC-2A contains four DC-41 controllers. The DC-41 is a dual-ported interface with only one port active at a time. Two DCC-2A cabinets are used in a DS-41R disk subsystem to provide dual-channel access.

Storage capacity and transfer rates are the same for both DS-41D and DS-41R disk subsystems.

DD-41 Features

Sustained transfer rate 9.6 Mbytes/s

Total data sectors 1,175,760

Total data words 601,989,120

Typical position delays:

Single track 5 ms

Average 16 ms

Full stroke 30 ms

DE-41 Power and Cooling

(four DD-41s)

Required power 208 Vac, 3 phase,
60 Hz, 20 A

Cooling air cooled

Heat load 8,000 Btu/hr, 2,340 W

DE-41 Physical Description

(four DD-41s)

Floor space 7.3 ft² (0.68 m²)

Weight 1,150 lbs (522 kg)

Cabinet dimensions:

Height 67 in. (170 cm)

Width 26 in. (66 cm)

Depth 41 in. (104 cm)

DD-41 Placement and Cabling

Minimum clearance:

Sides 1 in. (2.5 cm)

Front 36 in. (91.4 cm)

Back 30 in. (76.2 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of data cables ... 20 ft (6 m)

Four disk drives are connected to the DCC-2A chassis for the DS-41 disk subsystem. Eight disk drives are connected to the DCC-2A chassis for the DS-41D disk subsystem, but only four disk drives can be active at one time. This technique, known as daisy chaining, is used to double the capacity of a single subsystem from 19.2 Gbytes to 38.4 Gbytes; daisy chaining does not increase the 9.6-Mbyte/s transfer rate.

DD-49 Disk Drive Specifications



CRAY
RESEARCH, INC.

DD-49 Specifications

DD-49 Features

Storage capacity 1,200 Mbytes

Transfer rate:

Sustained 9.6 Mbytes/s

Burst rate 12 Mbytes/s

Total data sectors 297,696

Total data words 150,420,352

Typical position delays:

Single track 2 ms

Average 16 ms

Full stroke 30 ms

Power and Cooling

Required power 3 phase, 208 Vac,
60 Hz, 20 A per phase

Heat load 9,000 Btu/hr (2,640 W)

Type of cooling air cooled

Physical Description

Floor space 7.3 ft² (0.68 m²)

Weight 844 lbs (383 kg)

Placement and Cabling

Minimum clearance

Sides 12 in. (25 cm)

Front 36 in. (91 cm)


Back 30 in. (76 cm)

Length of power cable 6 ft (1.8 m)

Maximum length of

data cables 50 ft (15 m)

Front-end Interface Specifications



FEI Specifications

FEI Features

Cray Research, Inc. offers hardware interfaces and station software to connect the CRAY Y-MP C90 system to a wide variety of popular computer systems, networks, and workstations.

Mainframes:

- Amdahl 470 series
- CDC 70
- CDC 170
- CDC 180
- CDC 6000
- CDC 7600
- Honeywell 6000
- IBM 360
- IBM 370
- IBM 303x
- IBM 308x
- IBM 43xx
- Siemens
- Unisys 1100/80 series

Minicomputers and microcomputers:

- Data General ECLIPSE series
- DEC PDP/11
- DEC VAX 11/750
- DEC VAX 11/780
- DEC VAX 11/782
- DEC VAX 11/785
- DEC VAX 8600
- DEC VAX cluster
- Motorola Delta Series microcomputer

Networks:

- Ethernet (TCP/IP) networks
- Network Systems Corporation
- HYPERchannel

Workstations:

- Sun-3 (through FEI-3's interface)

Operating systems:

- Apollo AEGIS
- CDC NOS, NOS/BE, and NOS/VE
- Data General AOS
- Data General RDOS
- DEC VAX/VMS
- IBM MVS and VM
- Unisys
- UNIX

Physical Description

Floor space 4.38 ft² (3.42 m²)
Weight 200 lbs (91 kg)
Height 23 in. (58.4 cm)

FOL-3 Fiber-optic Link Specifications



FOL-3 Specifications

FOL-3 Description

The FOL-3 is a fiber-optic connection between a Cray Research I/O subsystem (IOS) and a front-end interface (FEI). The FOL-3 is an alternative to the wire cabling between an IOS and an FEI. The FOL-3 is designed primarily to increase the maximum cabling distance between a Cray Research computer system and a front-end computer and to provide complete electrical isolation from electromagnetic fields.

The FO cabinet is positioned on top of the IO cabinet. The FO cabinet contains an FO module that includes the receivers and transmitters for the fiber-optic cable and power connection for the module.

The IO cabinet contains a power supply and an IO module. The IO module provides an interface between the fiber-optic receiver/transmitter board and the Cray Research 6-Mbyte/s channel.

FOL-3 Configurations

The FOL-3 consists of the following equipment:


- Fiber-optic (FO) cabinet
- Interface (IO) cabinet
- Electrical kit

Below is an illustration of a general configuration of the FOL-3 used with an IOS. The dashed line encircles the components that compose the FOL-3.

At the Cray Research mainframe end (local end of the FOL) is a fiber-optic cabinet that consists of an IO cabinet and an FO cabinet.

At the FEI mainframe end (remote end of the FOL) of this link is an FEI cabinet. This cabinet consists of an FO cabinet positioned on top of an FEI cabinet. The FEI cabinet contains the modules necessary to communicate with the front-end computer system and a Cray Research IO module. The FO cabinet is identical to the FO cabinet at the local end of the link.

The electrical kit contains a Cray Research IO module, logic and power interconnections for the IO module, and logic and power interconnections for the signal connection to the FO module in the FO cabinet.



General FOL-3 Configuration

A-10398

FOL-3 Specifications

The following additional fiber-optic cable configurations are possible for the FOL-3:

- 3-Mbyte/s, 4-km cable
- 6-Mbyte/s, 2-km cable

The equipment configurations for 2-km and 4-km cable lengths are identical to those for the FOL-3.

The customer is responsible for supplying and installing the fiber-optic cables. A variety of cable types and vendors exists. See your local Cray Research sales representative for cable specifications.

FOL-3 Advantages

The following items are advantages of using the FOL-3 as opposed to wire cabling:

- Decreased cost
- Increased security
- Increased cabling distances
- Decreased vulnerability to interference
- Ease of handling

FOL-3 Features

The table shown below describes FOL-3 features.

Physical Description

Floor space 4.1 ft² (0.38 m²)
Weight 240 lbs (109 kg)
Height 27 in. (69 cm)

| FOL-3 Features | |
|--------------------------|---|
| Feature | Description |
| Fiber-optic Cable length | 3 ft. (0.91 m) to 3,280 ft. (1,000 m) |
| Power Requirements | -5.2 V, -2.0 V, 100-W total power |
| Transfer Rate | 3 Mbytes/s |
| Data Protection | Cyclic redundancy check (CRC) on link data, parity generation, and channel data check |
| Ground Isolation | Complete ground isolation between a Cray Research computer system and a front-end computer. |

6 SOFTWARE OVERVIEW

CRAY C90 series computer systems come with a variety of software, including the Cray Research operating system UNICOS. The CF77 compiling system provides automatic vectorizing, as do the Cray Research Standard C and Pascal compilers. Extensive library routines, program- and file-management utilities, debugging aids, a powerful Cray Research assembly language, and extensive support for industry standards are included in the system software. A large number of third-party and public-domain application programs also run on Cray Research systems.

CRAY C90 series computer systems are supported by industry standard communications software such as the International Standards Organization/Open Systems Interconnect (ISO/OSI) protocol and Transmission Control Protocol/Internet Protocol (TCP/IP). The CRAY C90 series systems are also supported by Cray Research proprietary station products for connecting to other vendors' systems and workstations.

UNICOS Operating System

CRAY C90 series computer systems come with the UNICOS operating system. The UNICOS operating system is derived from the UNIX Laboratories, Inc. UNIX System V operating system. It is also based in part on the Fourth Berkeley Software Distribution (BSD), under license from The Regents of the University of California.

The UNICOS operating system provides exceptional problem-solving ease; it provides powerful interactive and batch capabilities and multiple methods to accomplish a task. It efficiently manages high-speed data transfers between the CRAY C90 series system and peripheral equipment. The UNICOS operating system is written in C, a high-level language, and is available on all Cray Research systems.

The UNICOS operating system consists of a kernel plus a large set of utilities and library programs. The kernel is a simple structure with short and efficient software control paths. The kernel supports many system call primitives that library and application programs can use together to perform complex tasks.

The UNICOS operating system offers a large set of utility programs that allows the user to interact with the operating system. In addition, it provides a number of products specifically designed for Cray Research computer systems. The UNICOS operating system supports the following compilers: Fortran, Pascal, C, Lisp, and Ada.

The UNICOS operating system and UNIX are essentially the same in philosophy, structure, and function. However, Cray Research has enhanced UNIX to create the UNICOS operating system to use the power of the Cray Research computer system more efficiently. Enhancements include I/O capabilities to take advantage of supercomputer performance, added multiprocessor and multitasking support, additional networking software, accounting features, and others. The UNICOS operating system is designed for both interactive and batch environments. It supports the Network Queuing System (NQS) for batch processing.

Multiprocessing

Multiprocessing divides an application program into independent tasks called partitions and runs them in parallel. Compared to serially executed programs, multiprocessing can substantially improve throughput. Three multiprocessing features have evolved, and they can all work together in a single program, but not in the same program unit. The following subsections describe the three multiprocessing features.

Macrotasking Feature

Macrotasking is the first phase in the evolution of Cray Research parallel processing software. It requires extensive data scoping and insertion of Cray Research-specific library calls that allow parallel execution of code at the subroutine level on multiple processors. Macrotasking is best suited to programs with large, long-running tasks. The user interface to the system's macrotasking capability is a set of Fortran-callable subroutines that explicitly define and synchronize tasks at the subroutine level. These subroutines are compatible with similar subroutines available on other Cray Research products.

Microtasking Feature

Microtasking is the second phase in the evolution of Cray Research parallel processing software. Microtasking expands on the strengths of macrotasking but requires less data scoping and uses compiler directories rather than Cray Research-specific library calls.

Microtasking is a multiprocessing technique that allows parallel execution of very small segments of code on multiple processors. An example of this is individual iterations of DO loops. With microtasking, the programmer can revise the code or issue compiler directives to further enhance performance beyond the automatic vectorization done by the compiler.

In addition to working efficiently on parts of programs where the granularity is small, microtasking works well when the number of processors available for the job is unknown or may vary during the program's execution. Additionally, in a batch environment where processors may become available for short periods, the microtasked job can dynamically adjust to the number of available processors.

Autotasking Feature

The Autotasking feature (Autotasking) of the CF77 Fortran compiling system is the third phase in the evolution of Cray Research parallel processing software. Autotasking is based on the microtasking design and shares several advantages with microtasking: very low overhead synchronization cost, excellent dynamic performance independent of the number of central processing units (CPUs) available, both large and small granularity parallelism, and so on.

Autotasking has two fundamental improvements over microtasking. First, it is automatic multiprocessing. Autotasking allows user programs to be automatically partitioned over multiple CPUs (without user intervention). Second, Autotasking can exploit parallelism at the DO-loop level without extending to subroutine boundaries.

CF77 Compiling System

CRAY C90 series computer systems use the Cray Research CF77 compiling system. This compiling system is fully compliant with the ANSI 78 (Fortran 77) standards and offers a high degree of automatic scalar and vector optimization. The CF77 compiling system permits maximum portability of programs between different Cray Research systems and accepts many nonstandard programming constructs written for other vendors' compilers. Vectorized object code is produced from standard Fortran code; users can program in standard syntax to access the full power of the mainframe architecture.

The CF77 compiling system consists of the following software: the Autotasking Fortran Dependence analyzer, the Fortran translator, and the Cray Research Fortran 77 compiler (CFT77). This system is a multipass, optimizing, transportable compiling system that processes existing

standard Fortran programs. It uses two basic techniques to improve the execution time of a Fortran program: vectorization and scalar optimization.

The compiling system automatically generates code that uses the vector registers and functional units of the mainframe. The programmer does not need to know the details of vectorization because the compiling system automatically vectorizes Fortran programs. When the compiling system cannot vectorize code, it generates scalar code using a variety of optimization techniques to improve execution time. Scalar optimization transforms the internal representation of the Fortran program into a more efficient but functionally equivalent program.

The CF77 compiling system is portable on several levels. Because it is in compliance with the ANSI 78 standard, programs written for other computer systems have maximal portability to a CRAY C90 series system with minimal effort. Also, the compiling system is designed to run on all Cray Research systems, enabling a Fortran program that compiles and runs on one Cray Research system to run on all Cray Research systems. In general, programs that compile and execute correctly with the CFT compiler also compile and execute correctly with the CF77 compiling system.

C Compiler

The C language is a high-level system programming language. Most of the UNICOS kernel code and utilities are written in C because C is a structured and highly efficient language. Many programming applications are also written in C. The C language offers a large standard library of functions and an ever-expanding base of software application programs. The availability of C complements the scientific orientation of Fortran. The Cray Research Standard C compiler performs scalar optimization and vectorizes code automatically.

The Cray Research Standard C compiler is available on all Cray Research computer systems running the UNICOS operating system. The compiler translates C language statements into assembler instructions that make effective use of the Cray Research computer system.

The C preprocessor (cpp) is included as a part of the Cray Research Standard C compiler. The cpp enables macro substitution, conditional compilation, and the inclusion of named files in the compilation process.

Cray Research Standard C is portable on several levels. Because Cray Research Standard C is in compliance with the 1989 ANSI standard, programs written for other computer systems have maximal portability to a CRAY C90 series system with minimal effort.

Pascal

Pascal is a high-level, general-purpose programming language used as the implementation language for the CF77 compiling system and other Cray Research products. Cray Research Pascal complies with the ISO Level 1 standard and offers such extensions to the standard as separate compilation of modules, imported and exported variables, and an array syntax.

The Pascal compiler transforms Pascal code into machine language instructions that execute on Cray Research computer systems. Using Pascal, a programmer can implement algorithms and data structures in a high-level, machine-independent manner without sacrificing efficiency.

The Cray Research Pascal compiler takes advantage of the mainframe hardware features through scalar optimization and automatic vectorization. The compiler provides access to Fortran common block variables and uses a common calling sequence that allows Pascal code to call Fortran and CAL routines.

Cray Assembler

The Cray Assembly Language (CAL) enables a user to closely tailor a program to the architecture of the mainframe. Through CAL, a programmer may symbolically express all hardware functions of the mainframe. CAL allows the production of highly efficient machine language programs. The user may designate program and data information to enable complete control of the mainframe CPUs. This facilitates full use of various features, such as the shared text feature, whereby a single set of instructions can service many users simultaneously.

A set of versatile pseudo-operations for defining macro instructions and controlling the assembler enhances the basic instruction set. A macro library provides macro instructions for subroutine entry and exit, allowing for easy subroutine linkage.

Cray Ada Environment

The Cray Ada Environment includes a Cray Ada compiler and a set of related tools that link, debug, and maintain Ada application programs. The Cray Ada Environment also supports an implementation of Ada program libraries, providing for flexible, project-oriented software development. The Cray Ada Environment is validated under the current Ada Compiler Validation Capability (ACVC) test suite.

Cray Allegro CL

Cray Allegro CL is a complete implementation of Common Lisp. The Cray Allegro CL system consists of an interpreter, an optimizing compiler, and a set of functions. There are a number of extensions to the specification in Cray Allegro CL. Included among the extensions are the top level; the debugger; a foreign function interface; and Flavors, an object-oriented programming system. Cray Allegro CL was designed to be compact, fast, and robust with respect to detecting user errors. The implementation itself is written mostly in Common Lisp, with some portions written in the C language.

Subroutine Libraries

Cray Research software includes subroutines that are callable from the CF77 compiling system, C, Pascal, and CAL. The subroutines are divided into libraries, generally on a functional basis. Libraries containing various utilities, high performance I/O subroutines, and numerous math and scientific routines are available, as are special-purpose libraries.

Utilities

A broad variety of software tools assists both interactive and batch users in the efficient use of a CRAY C90 series computer system.

The SEGLDR segment loader is an automatic loader for code produced by the language processors CFT77, CFT, C, Pascal, and CAL that can also be explicitly controlled by the programmer. Program segments are loaded as required without explicit calls to an overlay manager.

The Cray Symbolic Debugger (CDBX) allows users to interactively detect program errors by examining both running programs and program memory dumps. Other debugging tools are available for dump analysis and interpretation.

A variety of performance aids assists in analyzing program performance and optimizing programs with minimal effort. These aids include both static and dynamic analyzers, as well as profilers for CPU and I/O usage. Many provide graphic user interfaces using the X Window System.

The UNICOS Source Manager utility tracks modifications to files. This system is useful when programs and documentation undergo frequent changes because of development, maintenance, or enhancement. Line- and screen-oriented text editors, such as vi and Emacs, offer versatility for users who wish to create and maintain text files. Other system utilities provide for proper management of the system resources.

Communications Software

A CRAY C90 series computer system fits into environments consisting of single or multiple Cray Research systems, other vendors' mainframes, minicomputers, workstations, and devices capable of high-speed data transfer. Cray Research provides easy user access to Cray Research system capabilities, the ability to distribute applications between Cray Research computer systems and other vendor systems, and effective integration into existing customer networks.

Communications and connectivity are supported by the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, the International Standards Organization/Open Systems Interconnect (ISO/OSI) protocol, and the UNICOS station call processor (USCP) protocol.

The TCP/IP product allows the CRAY C90 series computer system to function as a peer in TCP/IP-supported, open networking environments. TCP/IP is a set of computer networking protocols that enables two or more hosts to communicate. Further, it is a set of procedures that allows communication among all hosts on a network whether the systems are similar or not.

The TCP/IP networking protocols were defined by the U.S. Department of Defense and enhanced by the University of California at Berkeley with the UNIX system. TCP/IP is supported only under the UNICOS operating system.

The ISO/OSI protocol logically connects Cray Research systems to other systems running ISO/OSI protocols. The UNICOS operating system supports the File Transfers, Access, and Management (FTAM) and Virtual Terminal (VT) applications of OSI. FTAM provides an interactive file transfer service for unstructured text files, unstructured binary files, and file directory files. VT allows users to connect to a remote system from a Cray Research system and to use the resources of the remote system.

USCP provides, by way of station software products, support for communicating with various vendor systems through a vendor's proprietary networking capability, such as IBM's SNA or DEC's DECnet.

Cray Research station software products provide system access to proprietary protocol implementations through network gateways. Cray Research supplies the station software packages for various front-end systems. These packages support batch job submission, job status, job control, file transfer, and interactive access to Cray Research systems.

The following stations are available:

- Apollo station - provides the software connection between the Cray Research mainframe and the Domain workstation.
- CYBER station - joins the Cray Research system to the Control Data Corporation CYBER 180 series, 70/170, or 700/800 systems to form a powerful computing combination.
- VAX or VMS station - controls the hardware and software link between a DEC VAX computer system and a Cray Research computer system.
- MVS station - provides the software connection between an IBM System/370, Extended Architecture, or compatible computer system and a Cray Research computer system.
- VM station - enables IBM compatible systems running under control of the Virtual Machine/System Product (VM/SP) and Conversational Monitor System (CMS) to be linked with a Cray Research computer system.
- UNIX station - provides Cray Research operating system access to installations whose front ends run UNIX.
- SUPERLINK/MVS product - provides data access, application-to-application communication, and job processing between the UNICOS operating system and MVS systems.
- CLS-UX product - provides Cray Research operating system access to UNIX users through a VAX or VMS system.

Applications

Cray Research supports application software vendors in converting and optimizing software for CRAY C90 series computer systems. Many of the most widely used application programs are currently available and supported to run in the Cray Research UNICOS environment. These codes are in fields such as computational fluid dynamics, structural analysis, mechanical engineering, nuclear safety, circuit design, seismic processing, image processing, molecular modeling, and artificial intelligence.

Cray Research has also developed the UniChem and MPGS applications. UniChem is an integrated software environment for chemists. MPGS is an interactive postprocessing visualization tool.

The availability of applications for Cray Research systems is driven largely by customer requirements that are communicated to the software vendors. Cray Research supports the on-going process of converting and maintaining application software.

Software Publications

The following subsections provide a partial list of Cray Research software publications. The manuals provide additional information about the software described in this section. These manuals and other user publications can be ordered through Cray Research local or regional sales offices. Refer to the *User Publication Catalog* (publication number CP-0099) for a complete list of software publications.

UNICOS Operating System

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|--|
| SG-2005 | I/O Subsystem (IOS) Operator's Guide for UNICOS |
| SG-2017 | UNICOS Source Code Control System (SCCS) User's Guide |
| SG-2050 | UNICOS Text Editors Primer |
| SG-2052 | UNICOS Overview for Users |
| SG-2112 | UNICOS Installation Guide |
| SG-2113 | UNICOS System Administration |
| SR-2011 | UNICOS User Commands Reference Manual |
| SR-2012 | Volume 4: UNICOS System Calls Reference Manual |
| SR-2014 | UNICOS File Formats and Special Files Reference Manual |
| SR-2022 | UNICOS Administrator Commands Reference Manual |

Fortran

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|---|
| SR-3071 | CF77 Compiling System, Volume 1: Fortran Reference Manual |

C

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|---|
| SR-2074 | Cray Standard C Programmer's Reference Manual |

Pascal

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|-------------------------|
| SR-0060 | Pascal Reference Manual |

Libraries

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|---|
| SR-2057 | Volume 5: UNICOS Network Library Reference Manual |
| SR-2079 | Volume 1: UNICOS Fortran Library Reference Manual |
| SR-2080 | Volume 2: UNICOS Standard C Library Reference Manual |
| SR-2081 | Volume 3: UNICOS Math and Scientific Library Reference Manual |

Utilities

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|--|
| SD-2107 | I/O Subsystem Model E (IOS-E) Guide |
| SG-2051 | UNICOS Tape Subsystem User's Guide |
| SG-2094 | UNICOS CDBX Debugger User's Guide |
| SG-3074 | CF77 Compiling System, Volume 4: Parallel Processing Guide |
| SG-3078 | OWS-E Operator Workstation Operator's Guide |
| SG-3079 | OWS-E Operator Workstation Administrator's Guide |
| SR-0010 | Software Tools Reference Manual |
| SR-0066 | Segment Loader (SEGLDR) and ld Reference Manual |
| SR-2091 | UNICOS CDBX Symbolic Debugger Reference Manual |
| SR-3077 | OWS-E Operator Workstation Reference Manual |

Communications Software

| <u>Publication Number</u> | <u>Title</u> |
|-------------------------------|---|
| SA-0250 | Apollo DOMAIN Station Reference Manual |
| SC-0270 | CDC NOS/VE Link Software Command Reference Manual |
| SG-2009 | TCP/IP and OSI Network User's Guide |
| SI-0038 | IBM MVS Station Reference Manual |
| SI-0160 | IBM VM Station Command and Reference for COS |
| SI-0191 | SUPERLINK Guides |

Publication

| <u>Number</u> | <u>Title</u> |
|---------------|--------------|
|---------------|--------------|

| | |
|---------|--------------------------------------|
| SR-0034 | CDC NOS/BE Station Reference Manual |
| SR-0035 | CDC NOS Station Reference Manual |
| SU-0107 | UNIX Station User's Guide |
| SU-3121 | CLS-UX Station User Guide |
| SV-0020 | DEC VAX/VMS Station Reference Manual |

Applications

Publication

| <u>Number</u> | <u>Title</u> |
|---------------|--------------|
|---------------|--------------|

| | |
|------------|--|
| MCDR-1000N | Directory of Applications Software for Cray Research Supercomputers for 1993 |
|------------|--|

Software Training

Cray Research offers complete training on the software available for CRAY C90 series computer systems. Extensive user-support analyst and system analyst training is available at Cray Research's training facility. End-user and operator training are available at customer sites after installation of a Cray Research computer system. More information regarding courses and schedules can be obtained through your local or regional Cray Research sales office.

BIBLIOGRAPHY

Related Cray Research, Inc. hardware manuals are listed below. Refer to Section 6 for a list of related software manuals. To obtain Cray Research publications, order them from the Distribution Center:

Cray Research, Inc.
Distribution Center
2360 Pilot Knob Road
Mendota Heights, MN 55120
USA
800-284-2729 extention 35907

Cray Research Peripheral Equipment Site Planning Reference Manual,
CRI publication number HR-00080.

This manual provides site planning information about operator and maintenance workstation equipment, disk storage units (DSUs), and front-end interface (FEI) cabinets.

Cray Support Equipment Site Planning Reference Manual, CRI
publication number HR-00082.

This manual provides site planning information about refrigeration condensing units (RCUs) and motor-generator sets (MGSs).

CRAY Y-MP C90 Site Planning Reference Manual, CRI publication
number HR-04025.

This manual describes the physical requirements for the CRAY C90 series computer systems. It defines customer and Cray Research, Inc. site planning and preparation responsibilities. This manual also describes the operational requirements, system configurations, mainframe and cooling unit specifications and requirements, and computer room floor specifications.

Principles of Computer Room Design, CRI publication number HR-04013.

This manual describes computer room design principles to help computer room facility managers prepare, inspect, and maintain a stable, problem-free environment. Information on computer room and raised-floor construction, system cooling, environmental control, fire and lightning protection, power, and grounding is also discussed.

Safe Use and Handling of Fluorinert Liquids, CRI publication number HR-0306.

This manual is written for Cray Research, Inc. customers and field engineers whose Cray Research computer system uses Fluorinert Liquid. The manual warns and informs about using Fluorinert Liquid and describes its uses at Cray Research, Inc. The manual describes the Material Safety Data Sheet and explains its significance in using Fluorinert Liquid or any other chemical.

GLOSSARY

A

| | |
|--|--|
| A register | Address register. A registers are primarily used as address registers for memory references and as index registers. |
| Application | Software designed to perform a particular job or set of related jobs. |
| Autotasking | The process of automatically dividing up a program into individual tasks and organizing them to make the most efficient use of the Cray Research, Inc. computer hardware; a trademark of Cray Research, Inc. |
| Auxiliary input/output processor (EIOP) | A quarter board in the IOS-E that controls the transfer of data between the channel adapters (CAs) and the IOS-E buffer board. |

B

| | |
|-------------------|---|
| B register | Intermediate address register. B registers are used as intermediate storage for the A registers. |
| Bank | The smallest addressable division of central memory. |
| BDM | Bidirectional memory mode (bit). The modes field in the exchange package contains the BDM mode bit. When the BDM mode bit is set, block read and write operations can operate concurrently. |
| BiCMOS | Bipolar complementary metal oxide semiconductor. |
| BML | Bit matrix loaded (flag). The bit matrix loaded flag sets if the bit matrix has been successfully loaded. This bit in the exchange package is reloaded from memory on an exchange. |
| BMM | Bit matrix multiply functional unit. The BMM functional unit performs a logical multiplication of two matrices, designated A and B, creating a single-bit result for each pair of elements that is multiplied, which is designated matrix C. The result matrix C, is the product of matrix A and matrix B transposed (B^t). |
| BPI | Breakpoint interrupt (flag). The breakpoint interrupt flag sets if the interrupt-on-breakpoint (IBP) interrupt mode bit is set and enabled and a write reference is made to an address within the breakpoint range. |

B (continued)

Buffer board A module in the IOS-E that temporarily stores system data from the high-speed (HISP) channels or the channel adapters (CAs).

C

C90 mode The modes field in the exchange package contains the C90 mode bit. When the C90 mode bit is set, the mainframe operates in C90 mode. The V registers are 64 bits x 128 elements, the VL register is 8 bits wide, and the P register is 32 bits wide, which enables a program range of 1 Gword. The entire CRAY C90 series instruction set is also available.

CAL Cray Assembly Language. A symbolic language that generates machine instructions on a one-for-one basis and allows programs to call subroutines from the library through the use of pseudoinstructions.

CCA-1 A channel adapter that connects the IOS-E to a 6-Mbyte/s channel pair.

Central memory Memory residing in the mainframe.

Central processing unit (CPU) A module used in the mainframe that controls the flow of system data, performs mathematical and logical functions on system data, and executes program instructions.

Chaining The process of sequencing logical operations so the results of one operation may be used by another operation without needing a memory reference in between.

Channel adapter (CA) A component in the IOS-E that transfers control and data between the buffer board and the peripherals.

Checkbyte An 8-bit correction code (checkbyte) that is generated by the SECDED logic to protect each 64-bit word of data.

CLN Cluster number (register). The CLN register in the exchange package determines which set of the 17_{10} available clusters of SB, ST, and SM registers the CPU can access.

Cluster interface (CIN) A quarter board in the IOS-E that transfers control and data between the workstation interface (WIN) and the input/output processor multiplexer (IOP MUX) and auxiliary input/output processors (EIOPs).

Clusters A set of shared registers accessible by all CPUs. There are 17_{10} valid clusters of shared registers in a CPU.

C (continued)

| | |
|------------------------|---|
| Compiler | A software program used to convert high-level programming language into binary machine code. |
| CP | Clock period. The CP is the interval in which the system clock completes one oscillation. |
| CRAY C90 series | The CRAY C90 series consists of five product lines: the CRAY C92A, CRAY C94A, CRAY C94, CRAY C98, and CRAY C916 computer systems. |

D

| | |
|--------------|---|
| DA-60 | The Cray Research DD-60 high-performance disk array. |
| DA-62 | The Cray Research DD-62 high-performance disk array. |
| DBA | Data base address (register). The DBA register, part of the exchange package, holds the base address of the user's data range. |
| DCA-1 | A channel adapter that connects the IOS-E to DS-40, DS-41, and DS-49 disk storage units. |
| DCA-2 | A channel adapter that connects the IOS-E to DD-60, DD-61, and DD-62 disk drives. The DCA-2 disk channel adapter in the IOP manages control signals and protocol for the individual disk drives in a DE-60. |
| DCA-3 | A channel adapter that connects the IOS-E to a DD-60 or DD-62 disk array. The DCA-3 disk channel adapter in the IOP manages control signals and protocol for the DA-60 or DA-62. |
| DCC-2 | The Cray Research DCC-2 houses the DC-40, which is separate from the DD-40 disk drives. |
| DCU | Disk controller unit. An interface between the disk storage units (DSUs) and the auxiliary I/O processor (EIOP). |
| DC-40 | The Cray Research DC-40 disk controller provides interface logic to adapt DCA-1 signals and protocol for individual DS-40s, to handle routing among the drives, and to buffer data from the four spindles in a full-track buffer. |
| DC-41 | The Cray Research DC-41 disk controller provides interface logic to adapt DCA-1 signals and protocol for individual disk drive units, to handle routing among the drives, and to buffer data from the four spindles in a full-track buffer. |

D (continued)

| | |
|---------------------------|--|
| DD-40 | The Cray Research DD-40 disk drive. |
| DD-41 | The Cray Research DD-41 disk drive. |
| DD-49 | The Cray Research DD-49 disk drive. |
| DD-60 | The Cray Research DD-60 high-performance disk drive. |
| DD-61 | The Cray Research DD-61 high-performance disk drive. |
| DD-62 | The Cray Research DD-62 high-performance disk drive. |
| DE-60 | One Cray Research DE-60 disk enclosure cabinet contains a maximum of ten DD-60 and/or DD-62 disk drives. Eight of the disk drives store system data, and two disk drives are spares. The DCA-2 disk channel adapter in the IOP manages control signals and protocol for the individual disk drives in a DE-60. |
| Deadlock | A condition resulting in the inability to continue processing that is caused by an unresolvable conflict. A deadlock condition occurs when all CPUs in a cluster are holding issue on a test and set instruction. |
| Deadstart | The sequence of operations required to start an operating system running in a Cray Research computer system. |
| Dielectric coolant | A fluid that travels through the module cold plates, removes heat from the modules, and transfers the heat to refrigerant in the heat exchanger subassembly. |
| Disk array | A five-spindle disk array composed of DD-60 or DD-62 spindles supported by the DCA-3 channel adapter. Four of the spindles hold data, and the fifth spindle contains parity information on the data. The spindles are housed in DE-60 disk enclosure cabinets. |
| DL | Deadlock interrupt (flag). The deadlock interrupt flag sets if the IDL interrupt mode bit is set, the program is not in monitor mode, and a deadlock condition occurs because all CPUs in a cluster are holding issue on a test and set instruction. |
| DLA | Data limit address (register). The DLA register holds the upper limit address of the user's data range. |
| DRAM | Dynamic random-access memory. A memory device that must be refreshed periodically in order to store data. |
| DSU | Disk storage units. A computer disk drive. |

D (continued)

- DS-40** The DS-40 disk subsystem consists of the DD-40 disk drive, the DC-40 disk control unit (DCU), and the disk controller cabinet (DCC-2).
- DS-41** The DS-41 disk subsystem consists of the DC-41 disk controller and the DD-41 disk drive.

E

- EASE** An error acquisition software program (EASE) that records errors received through mainframe, IOS, and SSD error channels. EASE displays logged errors in an understandable format.
- EIM** Enable interrupt modes (flag). The interrupt modes field in the exchange package contains the EIM flag. An exchange to monitor mode or non-monitor mode sets the EIM flag.
- EIOP** Auxiliary I/O processor. The EIOP controls the channel adapters that connect the IOS-E to peripheral devices such as disk drives, tape drives, and communications channels.
- EMI** Electromagnetic interference. EMI is radiated energy that interferes with and distorts digital signals.
- ESL** Enable second vector logical mode (bit). The modes field in the exchange package contains the ESL mode bit. When the ESL mode bit is set, the second vector logical functional unit is enabled, and if it is not busy, it has first priority to execute instructions 140ijk through 145ijk.
- Ethernet** A particular type of network hardware that forms a physical link between computers; a trademark of Xerox Corporation.
- Exchange mechanism** The technique used in the CRAY C90 series computer system for switching instruction execution from program to program. Refer to exchange package.
- Exchange package** A 16-word block of data in memory reserved for exchange packages. The exchange package contains the necessary registers and flags associated with a particular program. Each program has its own exchange package.
- EXX** Error exit interrupt (flag). The EXX interrupt flag sets if the enable-interrupt-on-error-exit (FEX) interrupt mode bit is set and enabled and an error exit instruction (000000) issues. Issuing an error exit instruction always causes an exchange, regardless of the state of FEX.

F

| | |
|---------------------------------|--|
| F | Floating-point (operation). When an F appears in front of a register designator in a symbolic machine instruction, the calculation is a floating-point operation. |
| FEI | Front-end interface. An interface that connects the CRAY C90 series computer I/O channels to channels of front-end computers. An FEI compensates for differences in channel widths, machine word size, electrical logic levels, and control signals. |
| Fetch sequence | A fetch sequence transfers a block of instructions from memory to an instruction buffer. |
| FEX | Enable-interrupt-on error exit Mode (bit). The interrupt modes field in the exchange package contains the FEX bit. When the FEX bit is set, it enables an interrupt when an error exit occurs. |
| Floating-point operation | A mathematical or logical operation on two or more real numbers. |
| Fluorinert Liquid | The dielectric coolant circulated through the module cold plates; a trademark of 3M. |
| FNX | Interrupt-on-normal exit mode (bit). The interrupt modes field in the exchange package contains the FNX bit. When the FNX bit is set, it enables the NEX interrupt flag to set if a normal exit instruction (004000) issues. Issuing a normal exit instruction always causes an exchange, regardless of the state of FNX. This mode is not affected by the EIM flag. |
| FOL-3 | Fiber-optic link. The Cray Research 3-Mbyte/s fiber-optic link allows an FEI to be separated from a Cray Research computer system by distances of up to 3,281 ft (1,000 m). The FOL-3 provides complete electrical separation of the connected devices. |
| FPE | Floating-point error (flag). The interrupt flags field in the exchange package contains the FPE flag. The FPE flag sets when a floating-point range error occurs in any of the floating-point functional units and the Interrupt-on-floating-point error (IFP) flag is set. |
| FPS | Floating-point error status (bit). The status field in the exchange package contains the FPS status bit. The floating-point status bit sets if a floating-point error occurred during the execution interval. |
| Functional unit | Circuitry designed to perform a particular mathematical or logical operation. |

G

- Gate array** An array of circuits contained in a single integrated circuit package; these circuits may be customized in their operation as a group.
- Gather/scatter** An operation that places data at various intervals in the available memory storage and then gathers the data back into its original organization.

H

- H** Half-precision floating-point (operation). When an H appears in front of a register designator in a symbolic machine instruction, the calculation is a half-precision floating-point operation.
- HCA-3** A channel adapter that connects the IOS-E to a HIPPI input channel.
- HCA-4** A channel adapter that connects the IOS-E to a HIPPI output channel.
- HCA-5** A channel adapter that connects an IOP and to an external device.
- HDA** Head disk assembly. A sealed assembly that contains the magnetic storage media (disk drive), read and write heads, and servo mechanism.
- HEU** Heat exchanger unit. Part of the cooling equipment for the CRAY C90 series mainframe. The HEU uses a refrigerant to cool the dielectric coolant that circulates through the mainframe.
- High Performance Parallel Interface (HIPPI)** A type of interface used to transfer control and data between Cray Research, Inc. channel adapters and peripherals.
- High-speed (HISP) channel** A channel that transfers system data between the IOS-E and the mainframe or between the IOS-E and the SSD-E.
- High-speed control multiplexer (HCM)** A quarter board in the IOS-E that controls the transfer of high-speed (HISP) channel information between the IOS-E and the SSD-E or mainframe.
- HYPERchannel** A trademark and product of Network Systems Corporation that provides an interface between a LOSP channel and other brands of computers.

I

- I** Reciprocal iteration (operation). When an I appears in front of a register designator in a symbolic machine instruction, the calculation is a reciprocal iteration operation.
- IBA** Instruction base address (register). The IBA register is in the exchange package. The IBA register holds the base address of the user's instruction range.
- IBP** Interrupt-on-breakpoint mode (bit). The interrupt modes field in the exchange package contains the IBP bit. When the IBP bit is set, it enables an interrupt if a breakpoint occurs.
- ICM** Interrupt-on-correctable memory error mode (bit). The interrupt modes field in the exchange package contains the ICM bit. When the ICM bit is set, it enables interrupts on correctable memory data errors while data is being read from memory.
- ICP** Interprocessor interrupt (flag). The interprocessor interrupt flag sets if the IIP interrupt mode bit is set and enabled and another CPU requests an interrupt of this CPU by issuing instruction 0014j1.
- IDL** Interrupt-on-deadlock mode (bit). The interrupt modes field in the exchange package contains the IDL bit. When the IDL bit is set, it enables an interrupt if a deadlock occurs while the program is not in monitor mode. IDL has no effect in monitor mode.
- IFP** Interrupt-on-floating-point error mode (bit). The interrupt modes field in the exchange package contains the IFP bit. When the IFP bit is set, it enables interrupts on floating-point errors.
- IIO** Interrupt-on-I/O mode (bit). The interrupt modes field in the exchange package contains the IIO bit. When the IIO bit is set, it enables an interrupt if SIE is set and this CPU is the lowest-numbered CPU with IIO=1 and EIM=1.
- IIP** Interrupt-on-interprocessor interrupt mode (bit). The interrupt modes field in the exchange package contains the IIP bit. When the IIP bit is set, it enables an interprocessor interrupt if requested by another CPU.
- ILA** Instruction limit address (register). The ILA register is in the exchange package. The ILA register holds the limit address of the user's instruction field.
- IMC** Interrupt-on-request from MCU mode (bit). The interrupt modes field in the exchange package contains the IMC bit. When the IMC bit is set, it enables interrupts from the MCU.

I (continued)

| | |
|---|--|
| IMI | Interrupt-on-monitor mode instruction mode (bit). The interrupt modes field in the exchange package contains the IMI bit. When the IMI bit is set, it enables an interrupt if a monitor mode instruction (001ijk; j≠0) issues while the program is not in monitor mode. IMI has no effect in monitor mode. |
| Input/output cluster (IOC) | A component of the IOS-E that contains one cluster interface (CIN), one input/output processor multiplexer (IOP MUX), four auxiliary input/output processors (EIOPs), two high-speed control multiplexers (HCMs), one buffer board, sixteen channel adapters (CAs), and one programmable interrupt (PINT). |
| Input/output processor multiplexer (IOP MUX) | A quarter board in the IOS-E that transfers information between the high-speed control multiplexers (HCMs), the cluster interface (CIN), and the auxiliary input/output processors (EIOPs). |
| Input/output subsystem model E (IOS-E) | A component of a CRAY C90 series computer system that transfers system data between the peripherals, SSD solid-state storage device model E (SSD-E), and the mainframe. |
| Intelligent peripheral interface - 2 (IPI-2) | An interface used to transfer control and system data between Cray Research, Inc. channel adapters and peripherals. |
| Interrupt modes field | The interrupt modes field in the exchange package contains user-selectable bits that dictate the execution of the program. |
| Instruction buffer | A set of registers in a CRAY C90 series mainframe used for temporary storage of instructions before issue. Each instruction buffer can hold 128 consecutive instruction parcels. |
| Instruction fetch | The process of loading program code from central memory to an instruction buffer. |
| Instruction set | A set of instructions that a particular computer can perform. |
| I/O buffer | A buffer used to provide temporary storage for data transferred between the mainframe and peripheral devices. |
| IOC | Refer to input/output cluster. |
| IOI | I/O interrupt (flag). The I/O interrupt flag sets if the SIE bit is set and this CPU is the lowest-numbered CPU with I/O interrupt mode set and enabled when a LOSP or VHISP channel completes a transfer. |
| IOP | I/O processor. An IOP is a fast, multipurpose computer capable of transferring data at extremely high rates. The IOS-E contains multiple IOPs. |

I (continued)

- IOR** Operand range error mode (bit). The interrupt modes field in the exchange package contains the IOR bit. When the IOR bit is set, it enables interrupts on operand address range errors.
- IOS** An input/output subsystem; a trademark of Cray Research, Inc.
- IPC** Interrupt-on-request from programmable clock mode (bit). The interrupt modes field in the exchange package contains the IPC bit. When the IPC bit is set, it enables an interrupt on a request from the programmable clock.
- IPI-2** An interface used to transfer control and system data between Cray Research, Inc. channel adapters and peripherals.
- IPR** Interrupt-on-program range error mode (bit). The interrupt modes field in the exchange package contains the IPR bit. When the IPR bit is set, it enables the PRE interrupt flag to set if a program range error occurs.
- IRP** Interrupt-on-register parity error mode (bit). The interrupt modes field in the exchange package contains the IRP bit. When the IRP bit is set, it enables an interrupt if a register parity error is detected while data is being read from a register.
- IRT** Interrupt-on-request from real-time clock mode (bit). The interrupt modes field in the exchange package contains the IRT bit. When the IRT bit is set, it enables an interrupt on a request from the real-time clock.
- Issue sequence** The issue sequence selects the instruction indicated by the program address (P) register, decodes it, determines whether the required registers or functional units are available, and if so, enables the CPU to execute the instruction.
- IUM** Interrupt-on-uncorrectable memory error mode (bit). The interrupt modes field in the exchange package contains the IUM bit. When the IUM bit is set, it enables interrupts on uncorrectable memory data errors.

L

- Library** A set of commonly used software routines that are available to programmers and to programs being compiled.
- LLRC** Length/longitudinal redundancy check. An error control system based on the arrangement of data in blocks according to some preset rule, the correctness of each character within the block being determined on the basis of the specific rule or set.

L (continued)

Low-speed (LOSP) channel Low-speed channel. The LOSP channel has a transfer rate of 6 or 20 Mbytes/s and enables an I/O cluster to communicate with a mainframe, front-end interface, or SSD-E.

M

Mainframe A component of a CRAY C90 series computer system that contains central memory and central processing units (CPUs).

Maintenance channel A channel that connects the MWS-E to the SSD-E.

Maintenance control unit interface (MCUI) A quarter board in the mainframe that receives an interrupt signal from the programmable interrupt (PINT) and interrupts the specified CPU in the mainframe.

Maintenance workstation model E (MWS-E) A component of a CRAY C90 series computer system that provides an intelligent and dedicated platform for performing offline and online tests, monitoring environmental conditions, and recording hardware errors.

MCU Maintenance control unit. The maintenance control unit for the CRAY C90 series computer system is the MWS-E.

MCU Maintenance control unit interrupt (flag). The MCU interrupt flag sets if the IMC interrupt mode bit is set and enabled and the MCU interrupt signal becomes active on I/O channel 40.

MEC Memory error correctable interrupt (flag). The memory error (correctable) flag sets if the ICM interrupt mode bit is set and a correctable memory error occurs while data is being read from memory.

MEU Memory error uncorrectable (flag). The memory error (uncorrectable) flag sets if IUM interrupt mode is set and enabled and an uncorrectable memory error occurs while data is being read from memory.

MFC Mainframe chassis.

MGS Motor-generator set. An MGS converts primary power from commercial power mains to the voltage and frequency used by the mainframe power supplies.

MII Monitor instruction interrupt (flag). The monitor instruction interrupt flag sets if the IMI interrupt mode bit is set and a monitor mode instruction (001ijk; j ≠ 0) issues while the program is not in monitor mode.

M (continued)

| | |
|-------------------------|---|
| MM | Monitor mode (bit). The modes field in the exchange package contains the MM bit. When the MM mode bit is set, it inhibits all interrupts except memory errors, normal exit, and error exit. The program can execute those instructions that are privileged to monitor mode. |
| Monitor mode | A condition in which a CPU inhibits all interrupts except those caused by memory errors, normal exit, or error exit instructions. |
| Multiprocessing | Several computer processes or jobs being computed at the same time. |
| Multiprogramming | The process of writing software to use the capabilities of a computer to process multiple jobs simultaneously. |
| Multitasking | The capability to run two or more parts, or tasks, of a single program in parallel on different CPUs within a mainframe. |

N

| | |
|------------|---|
| NEX | Normal exit interrupt (flag). The interrupt flags field in the exchange package contains the NEX flag. The normal exit flag sets if the interrupt-on-normal-exit (FNX) interrupt mode bit is set and enabled and a normal exit instruction (00400) issues. Issuing a normal exit instruction always causes an exchange, regardless of the state of FNX. |
|------------|---|

O

| | |
|---|---|
| Operating system | The major controlling software running in a computer that controls its overall operation. |
| Operator workstation model E (OWS-E) | A component of a CRAY C90 series computer system that Cray Research, Inc. analysts and customers use to monitor the computer system. |
| ORE | Operand range error (flag). The interrupt flags field in the exchange package contains the ORE flag. The ORE flag sets when a data reference is made outside the boundaries of the DBA and DLA registers and the interrupt-on-operand range error bit is set. |

P

| | |
|--------------------------------------|--|
| P | Population count (operation). When a P appears in front of a register designator in a symbolic machine instruction, the calculation is a population count operation. |
| P register | Program address register. The P register selects an instruction parcel from one of the instruction buffers. The contents of the P register are stored in the program address register field in the exchange package. The P register is 24 bits wide in Y-MP mode and 32 bits wide in C90 mode. |
| Parcel | A 16-bit portion of a word that is addressable for instruction execution but not for operand references. |
| Parity | Equivalence in the check bit of transmitted and received data. |
| Pascal | A high-level programming language. |
| PCI | Programmable clock interrupt (flag). The programmable clock interrupt flag sets if the IPC interrupt mode bit is set and enabled and the counter in the programmable clock equals 0. |
| Pipelining | An operation or instruction that begins before a previous operation or instruction finishes. Pipelining is accomplished using fully segmented hardware. |
| PN | Processor number. The PN field in an exchange package indicates which CPU executed the exchange sequence. |
| Port | A hardware or software access path to memory. |
| PRE | Program range error (flag). The interrupt flags field in the exchange package contains the PRE flag. The PRE flag sets when an instruction fetch is made outside the boundaries of the IBA and ILA registers. |
| Programmable clock | A 32-bit counter in each CPU that is used to generate interrupts at selectable intervals. |
| Programmable interrupt (PINT) | A quarter board in the IOS-E that enables any input/output cluster (IOC) in the IOS-E to interrupt any CPU in the mainframe. |
| Protocol | Software that defines the precise way in which data is transferred from one place to another. |
| PS | Program status (bit). The interrupt modes field in the exchange package contains the PS bit. The PS bit is set by the operating system to denote whether a CPU concurrently processing a program with another CPU is the master or slave in a multitasking situation. |

Q

- Q** Parity count (operation). When a Q appears in front of a register designator in a symbolic machine instruction, the calculation is a parity count operation.

R

- R** Rounded floating-point (operation). When an R appears in front of a register designator in a symbolic machine instruction, the calculation is a rounded floating-point operation.
- RAM** Random access memory. A memory device that retains the stored data as long as power is applied. When power is removed from the device, the stored information is lost.
- RCU** Refrigeration and condensing unit. The RCUs dissipate the heat transferred from the heat exchanger units (HEUs).
- RD-62** The Cray Research RD-62 removable high-performance disk drive.
- Reciprocal approximation** The mathematical process of approximating the value of a real number when divided into one ($1/n$).
- Refrigerant** A fluid that removes heat from dielectric coolant and transfers the heat to water or air.
- Register** A hardware storage location for one word, byte, or element of data.
- Remote Support** A system that provides a network connection to a remote location through a Telebit NetBlazer router and Microcom high-speed modem. The system allows support personnel to dial into the site, log on the MWS-E, run maintenance tools, and monitor the Cray Research computer system.
- RISC** Reduced instruction set computer.
- RPE** Register parity error (flag). When a word is written into a register, a set of parity bits is generated and stored with the data bits. This set of parity bits is compared to another set that is generated when the word is read out of the register. An error is indicated when the two sets do not match. Parity errors set the register parity error (RPE) flag in the exchange package.
- RT** Load real-time clock (instruction). The RT instruction loads the real-time clock register with the contents of a scalar (S) register.

R (continued)

-
- RTC** Real-time clock. The RTC is a 64-bit counter that advances one count each clock period.
- RTI** Real-time interrupt (flag). The real-time interrupt flag sets if the IRT interrupt mode bit is set and enabled and a real-time interrupt request is received.

S

-
- S register** Scalar register. The S registers are the source and destination registers for operands executing scalar arithmetic and logical instructions.
- SB** Shared address (register). The SB register is a shared register used for transferring address information from one CPU to another.
- Scalar** A single numerical value that represents a single aspect of a physical quantity.
- Section** A major addressable division of central memory that may be further divided into subsections and banks.
- Segmentation** An operation that is divided into a discrete number of sequential steps, or segments. Fully segmented hardware is designed to perform one segment of an operation during a single clock period (CP).
- SEI** Selected for external interrupts (flag). The interrupt modes field in the exchange package contains the SEI flag. When the SEI flag is set, this CPU is preferred for I/O interrupts.
- Semaphore** A 1-bit value stored in a register and used by programs to communicate the occurrence of an event.
- Shared registers** Registers that are available for more than one CPU to write to and read from.
- Single-byte correction/
double-byte detection
(SBCDBD)** A method of detecting whether one or more 4-bit bytes in a word has an incorrect value. If only one byte has an incorrect value, that byte can be changed back to the correct value.
- Single-error
correction/double-error
detection (SECDED)** A method of detecting whether one or more bits in a word has an incorrect value. Each 64-bit word of data is protected with an 8-bit correction code (checkbyte) generated by the SECDED logic. If only 1 bit in the 64-bit word has an incorrect value, that bit can be changed back to the correct value by the SECDED logic.

S (continued)

| | |
|---|--|
| SM | Semaphore register. The semaphore registers allow a CPU to temporarily suspend program operation in order to synchronize operation with other CPUs. |
| Small computer system interface (SCSI) devices | Components of the MWS-E and OWS-E that store and retrieve data used in the workstations. |
| Source code | A software program written in a high-level programming language. |
| SPARC | Scaleable Processor ARChitecture. A trademark of SPARC International, Inc. |
| Spindle | A component of a disk drive. |
| SSD solid-state storage device model E (SSD-E) | A component in a CRAY C90 series computer system that provides secondary data storage for the IOS-E and the mainframe; SSD is a trademark of Cray Research, Inc. |
| SSD-E/32i | A single coldplate component in a CRAY C92A or CRAY 94A computer system that provides 32 Mwords of secondary data storage for the IOS-E and the mainframe. |
| ST | Shared T (register). The ST register is a shared register used for transferring data from one CPU to another. |
| Status field | The exchange package contains a status field that is used to determine the operating modes of a CPU. |
| Subsection | A major addressable division of memory that can be further divided into banks. |
| System Maintenance and Remote Testing Environment (SMARTE) | An online program that performs hardware verification, error detection, error isolation, and automated degradation of faulty hardware components; SMARTE is a trademark of Cray Research, Inc. |

T

| | |
|-------------------|---|
| T register | Intermediate scalar register. The T registers are used as intermediate storage for the S registers. |
| TCA-1 | Tape subsystem channel adapter. The TCA-1 tape channel adapter transfers data between the I/O buffer and tape controllers. |
| TCA-2 | Tape subsystem channel adapter. The TCA-2 channel adapter provides an interface between an IOP and an external tape controller. |

U

- UNICOS** An operating system for Cray Research computer systems based primarily on the UNIX System Laboratories, Inc. UNIX System V and partially on the Fourth Berkeley Software Distribution. UNICOS is essentially the same in philosophy, structure, and function as UNIX, but has been enhanced to exploit the power of Cray Research computer systems. UNICOS is a trademark of Cray Research, Inc.
- UTC-1** Universal time clock channel adapter. The UTC-1 provides resident application programs with read access to the current Greenwich mean time and day-of-year.

V

- V register** Vector register. Each V register contains 64 bits x 64 elements in Y-MP mode and 64 bits x 128 elements in C90 mode.
- Vector** A single numerical value that contains information on more than one aspect of a physical quantity.
- Very high-speed (VHISP) channel** A channel that transfers system data between the mainframe and the SSD-E.
- VL** Vector length (register). The program-selectable VL register controls the effective length of a vector register for any operation. The VL register is 7 bits wide in Y-MP mode and 8 bits wide in C90 mode.
- VM** Vector mask (register). The VM field allows for the logical selection of particular elements of a vector.
- VNU** Vector not used status (bit). The state of the VNU bit in the exchange package status field indicates whether vector instructions (077xxx or 140xxx through 177xxx) were issued during the execution interval.

W

- Warning and control system (WACS)** A system that monitors the refrigeration and power distribution systems to ensure that the computer system is operating within recommended temperature and voltage limits.
- Word** A set amount of data that contains 64 bits of system data and 8 check bits.
- Workstation interface (WIN)** A quarter board in the IOS-E that transfers control and data between the MWS-E or OWS-E and the cluster interface (CIN).

W (continued)

- WS** Waiting for semaphore status (bit). The interrupt modes field in the exchange package contains the WS status bit. The waiting on semaphore bit sets if a test and set instruction (0034jk) is holding issue.

X

- XA** Exchange address (register). The XA register in the exchange package specifies the address of the first word of a 16-word exchange package loaded by an exchange sequence.

Y

- Y-MP mode** Y-MP mode is selected by setting the C90 mode bit in the exchange package to 0. When the mainframe is operating in Y-MP mode, the V registers are 64 bits x 64 elements, the VL register is 7 bits wide, and the P register is 24 bits wide, which enables a program range of 4 Mwords. In Y-MP mode, only instructions defined in previous CRAY Y-MP systems are available.

Z

- Z** Leading-zero count (operation). When a Z appears in front of a register designator in a symbolic machine instruction, the calculation is a leading-zero count operation.

INDEX

Boldface numbers refer to figures and tables.

Numbers

- 24-bit integer multiply (performed in a floating-point multiply functional unit), **2-15**
- 32-bit integer multiply (performed in a floating-point multiply functional unit), **2-16**

A

- A register
 - applications, 2-6
 - computation section, 2-4
 - fields, 2-37
 - transfer instructions, 2-57, 2-59
- Access conflicts, 2-43
- ACVC, 6-5
- Adapters, network, 5-14
 - See also* Channel: adapters
- Adding coefficients. *See* Algorithms
- Address add functional unit, 2-8
 - See also* Integer arithmetic
- Address data, 2-4
- Address multiply functional unit, 2-8
 - See also* Integer arithmetic
- Address registers. *See* A register
- Algorithms
 - floating-point addition, 2-20–2-21
 - floating-point division, 2-22–2-26
 - floating-point multiplication, 2-21–2-22
 - functional units, 2-7
- Alternate-path configuration
 - DD-60, **5-5**
 - DS-41, 5-10
- AND function, 2-13
- Apollo station, 6-8
- Application software, 6-8–6-9
- Approximating roots, 2-23
- Autotasking, 2-39, 2-47, 6-3

B

- B registers, 2-4, 2-6
 - transfer instructions, 2-61
- Biased and unbiased exponent ranges, **2-17–2-18**
- BiCMOS chips, 2-1
- Bidirectional memory transfer instructions, 2-63
- Bit count instructions, 2-74
- Bit matrix multiply
 - arithmetic, 2-27–2-29
 - instructions, 2-67
- Block diagrams
 - CRAY C90 series computer system, **1-3**
 - CRAY C916 cooling system, **1-17**
 - CRAY C92A and CRAY C94A cooling system, **1-11**
 - CRAY C94 and CRAY C98 cooling system, **1-14**
 - mainframe, **2-2**
- Block transfers, 2-6
- Branch instructions, 2-75–2-77
- Breakpoint interrupt instructions, 2-80

C

- C compiler, 6-4
- C90 mode and Y-MP mode differences, 2-49
- Cabinet, fiber-optic, 5-15
- Cabling. *See* Fiber-optic link
- CAL, 6-5
- CAL syntax, special forms. *See specific instructions*
- CCA-1 channel adapter, 3-4
- CDBX, 6-6
- Central memory
 - access conflicts, 2-43
 - as a functional unit, 2-8
 - general, 2-1–2-3

- Central memory (continued)
 - overview, 1-2
 - RAM, 2-1
- CF77 compiling system, 6-3-6-4
- Chaining
 - instruction sequence, 2-45
 - vector, 1-4
 - vector example, 2-44, 2-45
- Channel, 1-2
 - control instructions, 2-78
 - HISP, 1-2, 1-4, 3-1, 3-3-3-5, 4-3-4-4
 - LOSP, 1-2, 3-1, 3-3-3-5
 - VHISP, 1-2, 1-5, 4-3
- Channel adapters, 1-5, 3-1, 3-4-3-8
- Chilled water system
 - CRAY C916, 1-16-1-17, 1-17
 - CRAY C92A and CRAY C94A, 1-10, 1-11
 - CRAY C94 and CRAY C98, 1-13, 1-14
- CLN. *See* Cluster number
- CLS-UX product, 6-8
- Cluster 0 channel connections, 3-2
- Cluster number
 - field, 2-37
 - set instruction, 2-79
- Clusters
 - components of, 3-1
 - interprocessor communication section, 2-3
 - IOS-E, 1-5, 3-1-3-8
- Coefficients, adding. *See* Algorithms
- Common lisp, 6-6
- Communications software, 6-7-6-9
- Computation section, CPU, 2-4-2-28
- Concurrent operations, chaining, 2-45
 - See also* Multiprocessing; Functional unit:
independence
- Conditional branch instructions, 2-76
- Configurations
 - See also* Alternate-path; Daisy chain;
Single-port
 - CRAY C916 computer system, 1-28
 - CRAY C92A computer system, 1-20
 - CRAY C94 computer system, 1-24
 - CRAY C94A computer system, 1-22
 - CRAY C98 computer system, 1-26
 - DS-41A field-upgradable, 5-11
- Conflicts, central memory access, 2-43
- Control section, CPU, 2-30-2-38
- Conventions, instruction, 2-49-2-50
- Conversion, floating-point to decimal, 2-17
 - See also* Normalized floating-point numbers
- Cooling and support equipment
 - CRAY C916, 1-15-1-18
 - CRAY C92A and CRAY C94A, 1-9-1-18
 - CRAY C94 and CRAY C98, 1-12-1-18
- Cpp, 6-4
- CPU
 - computation section, 2-4-2-28
 - control section, 2-30-2-38
 - instruction summary, 2-55-2-80
 - overview, 1-1
 - programmable clock, 2-38, 2-79
 - shared resources, 1-2, 2-1-2-4
- Cray Ada Environment, 6-5
- Cray Allegro CL, 6-6
- Cray assembler, 6-5
- Cray Assembly Language, 6-5
- CRAY C90 series
 - components, 1-1
 - computer system block diagram, 1-3
 - disk storage units, 1-5
 - functional unit operations, 2-13-2-29
 - interrupt flags, 2-35
 - interrupt modes, 2-33
 - mainframe block diagram, 2-2
 - mainframe specifications, 2-81
 - maintenance and monitoring, 1-7
 - network interfaces, 5-14-5-17
 - operating modes, 2-36
 - software, 6-1
 - status field bit assignments, 2-36
- CRAY C916 computer system
 - configurations, 1-28
 - cooling system block diagram, 1-17
 - power and cooling equipment, 1-15-1-18
- CRAY C92A and CRAY C94A computer systems,
 - configurations, 1-20, 1-22
 - cooling system block diagram, 1-11
 - power and cooling equipment, 1-9-1-12
- CRAY C94 and CRAY C98 computer systems
 - configurations, 1-24, 1-26
 - cooling system block diagram, 1-14
 - power and cooling equipment, 1-12-1-14
- CRC errors, 3-5
- CYBER station, 6-8

D

Daisy chain configuration
 DD-60, 5-4
 DS-40, 5-13
 DS-41, 5-10

Data
 errors, 1-4
 flow (computation section), 2-4-2-5
 formats, 1-4
 storage in central memory, 2-1
 transfer, 1-2, 2-3, 2-6
 transfer registers, 2-6

Data base address register field, 2-32

Data limit address register field, 2-32

Data transfers
 from S and V registers, 2-44
 mainframe to SSD-E, 4-3-4-6
 mainframe to SSD-E/32i, 4-5
 SSD-E to IOS-E, 4-3-4-4
 SSD-E/32i to IOS-E, 4-6-4-7

DBA register, 2-32

DC-40 functions, 5-12

DC-41 disk controller, 5-9, 5-10

DCA-1
 channel adapter, 3-5
 functions, 5-9
 general, 1-5

DCA-2
 channel adapter, 3-5
 functions, 5-2, 5-6
 general, 1-5

DCA-3
 channel adapter, 3-5
 functions, 5-8
 general, 1-5

DCC-2, 5-11

DCU, 5-11

DD-49. *See* Disk drives

DD-60. *See* Disk drives

DD-61. *See* Disk drives

DD-62. *See* Disk drives

DE-60 cabinet, 5-2

DEC VAX supercomputer gateway, 5-17

Diagnostics, 1-7

Disk array
 block diagram, 5-8
 description, 5-8

Disk drives
 channel adapters, 3-5
 DD-49, 5-13
 DD-60, 5-1, 5-4, 5-5
 DD-61, 5-6-5-7
 DD-62, 5-7
 RD-62, 5-8
 specifications
 DA-60, 5-27-5-28
 DA-62, 5-29-5-30
 DD-49, 5-35-5-36
 DD-60, 5-19-5-20
 DD-61, 5-21-5-22
 DD-62, 5-23-5-24
 RD-62, 5-25-5-26

Disk storage units, 1-5

Disk subsystem specifications
 DS-40 and DS-40D, 5-31-5-32
 DS-41, DS-41D, and DS-41R, 5-33-5-34

Disk subsystems
 DS-40, 5-11-5-13
 DS-41, 5-9-5-11

Division
See also Algorithm; Floating-point; Reciprocal approximation
 algorithm, floating point, 2-22-2-26
 integer, 2-16

DLA register, 2-32

Documentation. *See* Publications

Double-precision numbers, 2-27

DRAM chips, 4-1

DS-41B package, 5-11

DS-40. *See* Disk subsystem specifications; Disk subsystems

DS-41. *See* Disk subsystem specifications; Disk subsystems

DS-41A field-upgradable configurations, 5-11

DSUs, 1-5

E

EIM flag, 2-33, 2-34

EIOPs
 general, 1-5
 I/O buffer, 3-3
 I/O cluster, 3-1, 3-2

Electrical separation, 1-9

- Elements, equalizing. *See* Algorithms
 - Equipment separation, 1-9
 - Equivalence, logical, instructions, 2-71
 - Error exit instructions, 2-77
 - Error logging programs, general, 1-7
 - Error messages. *See* Special: register values
 - Exchange
 - address (XA) register field, 2-37
 - address set instruction, 2-78
 - mechanism, 2-30
 - package fields, 2-31–2-37
 - sequence, 2-30
 - Exclusive NOR function, 2-13
 - Exclusive OR
 - function, 2-13
 - logical, instructions, 2-71
 - Exponent ranges, biased and unbiased, 2-17
 - Expression (*exp*), 2-55
- F**
- FEI, 1-8
 - FEI-1 front-end interface, 5-14–5-15
 - FEI-3 front-end interface, 5-15
 - specifications, 5-38–5-40
 - FEX mode, 2-32, 2-34
 - Fiber-optic link, 5-14–5-15
 - See also* FOL-3 specifications
 - Field-upgradable configurations, DS-41A, 5-11
 - Fields, exchange package, 2-31–2-37
 - Fixed-point operations. *See* Integer arithmetic
 - Floating-point
 - add and multiply range errors, 2-19
 - arithmetic, 2-5, 2-16–2-29
 - arithmetic instructions, 2-67–2-69
 - computations, 1-2
 - conversion to decimal, 2-17–2-18
 - data format, 2-16–2-17
 - functional unit underflow condition, 2-19–2-20
 - functional units, 2-11–2-12
 - numbers, normalized, 2-18
 - range errors, 2-19–2-20
 - Floating-point algorithms
 - addition, 2-20–2-21
 - division, 2-22–2-26
 - multiplication, 2-21–2-22
 - Floating-point multiply functional unit
 - division algorithm, 2-23–2-26
 - integer arithmetic, 2-14–2-16
 - normalized numbers, 2-12, 2-18
 - Floating-point reciprocal approximation
 - See also* Reciprocal approximation: functional unit
 - functional unit, 2-12
 - range errors, 2-20
 - Floating-point add functional unit, 2-12
 - normalized numbers, 2-12, 2-18
 - Floating-point reciprocal approximation, instructions, 2-69
 - Fluorinert Liquid, warning, 1-16
 - FNX mode, 2-32, 2-33
 - FOL-3 specifications, 5-40–5-43
 - See also* Fiber-optic link
 - Formats, instruction, 2-50–2-53
 - FPE, range errors, 2-19, 2-20
 - Front-end interface specifications, 5-38–5-39
 - See also* FEI
 - FTAM applications, 6-7
 - Functional instruction summary, 2-56
 - Functional unit operations
 - 24-bit integer multiply, 2-15
 - 32-bit integer multiply, 2-16
 - approximating roots, 2-24
 - biased and unbiased exponent ranges, 2-17
 - floating-point add and multiply range errors, 2-19
 - floating-point arithmetic, 2-16–2-29
 - floating-point data format, 2-16
 - floating-point reciprocal approximation range errors, 2-20
 - integer arithmetic, 2-14–2-16, 2-16–2-18
 - integer data formats, 2-14
 - internal representation of a floating-point number, 2-17
 - logical operations, 2-13–2-14
 - Functional units
 - address, 2-8
 - floating-point, 2-11–2-12
 - general, 1-2, 2-7–2-8
 - independence, 2-39, 2-42
 - instruction summary, 2-56
 - population/parity/leading 0 count, 2-8, 2-9

Functional units (continued)

- scalar, 2-8–2-9
- vector, 2-9–2-11

G

Graphics, raster, 5-16

H

Hardware. *See* Pipelining and segmentation

HCA-3 and HCA-4 channel adapters, 3-6, 5-16

HCA-5 channel adapter, 3-7

HEU

CRAY C916, 1-15, **1-17**

CRAY C94 and CRAY C98, 1-13, **1-14**

HIPPI

channel adapters, 3-6

external channel, 5-16

HISP channels

general, 1-2

I/O cluster, 3-1, 3-3–3-4

transfers, 4-3

HYPERchannel adapter, 5-16

I

I/O buffers, 3-3

I/O cluster. *See* Clusters

I/O section, 1-2, 2-3

IBA register, 2-31

IBM compatible magnetic tape drives and controllers, 1-6

IFP mode, range errors, 2-19, 2-20

ILA register, 2-32

Inclusive OR function, 2-13

Index calculation, 2-4, 2-6

Instruction

See also Functional units; Instructions

chaining sequence, **2-45**

differences between Y-MP mode and C90 mode, 2-50–2-53

execution, switching from program to program, 2-30

fetch sequence, 2-38

formats, 2-50–2-53

issue, 1-2, 2-38

pipelining and segmentation, 2-39

set, 2-5

summary, CPU, 2-55–2-80

Instruction base address register field, 2-31

Instruction limit address register field, 2-32

Instructions

bit count, 2-74–2-75

bit matrix multiply, 2-67–2-68

branch, 2-75–2-77

channel control, 2-78

error exit, 2-77

floating-point arithmetic, 2-67–2-69

functional unit, summary, 2-56

functional, summary, 2-56

integer arithmetic, 2-65–2-66

interprocessor interrupt, 2-80

interregister transfer, 2-59–2-62

logical operation, 2-69–2-72

memory transfer, 2-63–2-65

monitor mode, 2-54, 2-78–2-80

normal exit, 2-77

operand range error interrupt, 2-79

register entry, 2-57–2-59

shift, 2-73–2-74

V register transfer, 2-58, 2-61

vector, 2-44

vector population count, 2-74

Write, 2-64

Integer

computations, 1-2

data formats, **2-14**

product, 32-bit, 2-4

Integer arithmetic

address functional units, 2-8

general, 2-5, 2-14–2-16, **2-16**

instructions, 2-8–2-10, 2-65–2-67

scalar functional units, 2-8–2-10

Interfaces, network, 5-14–5-17

See also FEI

Intermediate registers, 2-5

address (B), 2-4

scalar (T), 2-4

transfer instructions, 2-61

Internal representation of a floating-point number, **2-17**

Interprocessor

communication section, 1-2, 2-3

interrupt instructions, 2-80

Interregister transfer instructions, 2-59–2-62

Interrupt
 flags field, 2-34–2-35, **2-35**
 modes field, 2-32–2-33, **2-33**

Interrupt-on-register-parity error mode. *See* IRP

IOP
 I/O cluster, 3-3–3-4
 instructions, 3-3
 MUX, 1-5, 3-1, 3-2

IOR mode, 2-32

IOS-E
 clusters, 1-5, 3-1–3-7
 data transfers, 4-3–4-4
 functions, 1-4–1-5
 specifications, 3-11–3-14

IPR mode, 2-32, 2-34

IRP, 2-6, 2-7

ISO/OSI protocol, 6-7

Iterations, based on Newton's method, **2-24**, 2-25

K

Kernel, UNICOS, 6-1

L

Libraries subroutine, 6-6

LLRC errors, 3-6

Logical
 operation instructions, 2-69–2-73
 operations, 2-13–2-14

LOSP channel, I/O cluster, 1-2, 3-1, 3-3–3-4

M

Macrotasking, 6-2

Mainframe
 block diagram, **2-2**
 data transfers, 4-3, 4-5
 overview, 1-2–1-4
 specifications, 2-81

Maintenance with MWS-E, 1-6

Maintenance workstation. *See* MWS-E

Matrix operations, indexing for, 2-4

Memory
 SSD-E, **4-2**
 transfer instructions, 2-63–2-65

Merge
 instructions, 2-72–2-73
 operation, 2-14

MGSs
 CRAY C916, 1-15–1-16
 CRAY C94 and CRAY C98, 1-12

Microtasking
 features, 6-2–6-3
 overview, 2-46–2-47

Modes field, **2-36**

Monitor mode instructions, 2-54, 2-78–2-80

Monitoring, with MWS-E, 1-7

Motor generator sets. *See* MGSs

MPGS interactive postprocessing visualization tool, 6-8

MPU card, 5-13

Multiplication. *See* Algorithms; Integer arithmetic

Multiprocessing
 defined, 1-1, 1-4, 2-39
 features, 6-2–6-3
 overview, 2-46–2-47

Multitasking, 1-4, 2-39, 2-46–2-47

MVS station, 6-8

MWS-E
 functions, 1-7–1-9
 WINs, 3-8
 workstation chassis, **1-6**

N

Network
See also FEI
 connections, direct, 5-16
 gateways, 6-7
 interfaces, 5-14–5-17

Networking protocol, 6-7

Newton's method for approximating roots, 2-23, **2-24**

No-operation instruction. *See* Monitor mode instructions

NOR function, exclusive, 2-13

Normal exit instructions, 2-77

Normalized floating-point numbers, 2-18

Normalizing results. *See* Algorithms

Notational convention for instructions, 2-49–2-50

O

- Offline diagnostic
 - listings, 1-7
 - testing, 1-7
- OpenWindows, general, 1-6
- Operand range error interrupt instructions, 2-79
- Operating
 - modes, 2-36
 - registers, 2-5-2-7
- OR function
 - exclusive, 2-13, 2-71
 - inclusive, 2-13
- ORE interrupt flag, 2-32
- Overflow condition. *See* Range errors, floating-point
- OWS-E
 - functions, 1-6, 1-8
 - WINs, 3-8
 - workstation chassis, 1-6

P

- P register, CPU control section, 2-30
- Parallel processing features, 2-39-2-47
- Parity bits, 2-6, 2-7
- Pascal, 6-5
- Performance counter instructions, 2-80
- Performance monitor, 2-38
- Peripheral devices. *See* Channel: adapters
- Pipe
 - 0 functional units, 2-9
 - 1 functional units, 2-9
- Pipelining and segmentation, 2-39-2-41
 - See also* Segmentation
- Pipes, vector operations, 2-7
- Population parity count instructions, 2-75
- Population/parity/leading zero count functional unit, 2-8, 2-9
- Ports, CPU, 2-1
- Power and cooling equipment
 - CRAY C916, 1-15-1-18
 - CRAY C92A and CRAY C94, 1-9-1-18
 - CRAY C94 and CRAY C98, 1-12-1-18
- PRE interrupt flag, 2-32
- Pressure monitoring. *See* WACS
- Primary registers, 2-5, 2-8
- Processor number field, 2-37

- Program address register. *See* P register, CPU control system
- Program address register field, 2-31
- Program range, 2-55
- Programmable clock
 - general, 2-38
 - instructions, 2-79
- Programmable real-time interrupt, 3-9
- Publications
 - software, 6-9-6-11
 - training, 6-11

R

- RAM, central memory, 2-1
- Range errors, floating-point, 2-19-2-20
- Raster graphics, 5-16
- RCU-5A, description, 1-16
- RCU-9, description, 1-13, 1-16
- Read instructions, 2-64-2-65
- RD-62. *See* Disk drives
- RDE-6 enclosure, 5-8
- Real-time clock
 - general, 2-4
 - set instruction, 2-78
- Reciprocal approximation
 - See also* Floating-point reciprocal approximation
 - floating-point, instructions, 2-69
 - functional unit
 - division algorithm, 2-22-2-26
 - normalized numbers, 2-18
- Refrigeration condensing unit. *See* RCU-5A, description; RCU-9, description
- Register
 - access, 2-5
 - entry instructions, 2-57-2-59
 - values, special, 2-54
- Register parity error flag. *See* RPE
- Registers
 - address, 2-6
 - intermediate and primary, 2-5
 - operating, 1-2, 2-5-2-7
 - scalar, 2-6-2-7
 - vector, 2-7
- Remote support, 1-7
- Reservations on V registers, 2-44
- Results, normalizing. *See* Algorithms

Return jump instructions, 2-77

RPE, 2-6, 2-7

S

S register

See also specific scalar functional units
field, 2-37

transfer instructions, 2-57, 2-60–2-61

SB registers, 2-3

SBCDBD, 2-1

Scalar

data, 2-4

floating-point data, 2-4

leading zero count instruction, 2-75

memory references, 2-6

population count instructions, 2-74

processing overview, 1-1, 1-4

register uses, 2-6–2-7

(S) registers, 2-4

segmentation and pipelining example, 2-40

Scalar add functional unit, 2-9

See also Integer arithmetic

Scalar instructions. *See Floating-point arithmetic*

Scalar logical functional unit, 2-9

Scalar shift functional unit, 2-9

SECDDED, 2-1, 3-3, 4-2

Segmentation

functional unit, 2-8

general, 2-39–2-43

scalar example, 2-40

vector example, 2-41

Semaphore. *See SM registers*

Separation

electrical, 1-9

equipment, 1-9

Shared register transfer instructions, 2-61

Shared registers, 2-3

Shared resources, CPU, 1-2, 2-1–2-4

See also Specifications

Shift instructions, 2-73–2-74

Single-port configuration

DD-60, 5-3

DS-40, 5-12

DS-41, 5-10

SM registers, 2-3

transfer instructions, 2-59

SMARTe platform, 1-8

Software

CRAY C90 series, 6-1

publications, 6-9, 6-11

Special

CAL syntax forms, 2-54

register values, 2-54

syntax with integer arithmetic instructions, 2-65

Specifications

DA-60, 5-27–5-28

DA-62, 5-29–5-30

DD-49, 5-35–5-36

DD-60, 5-19–5-20

DD-61, 5-21–5-22

DD-62, 5-23–5-24

DS-40 and DS-40D, 5-31–5-32

DS-41, DS-41D, and DS-41R, 5-33–5-34

FOL-3, 5-40–5-43

front-end interface, 5-38–5-39

IOS-E, 3-11–3-14

mainframe, 2-81

RD-62, 5-25–5-26

SSD-E, 4-7–4-9

SSD-E/32i, 4-9

SSD-E

data transfers, 4-3–4-6

LOSP channel testing, 1-8

memory, 4-2

memory sizes, 4-2

overview, 1-5

specifications, 4-7–4-9

SSD-E/32i

data transfers, 4-5–4-8

LOSP channel testing, 1-8

memory, 4-5

overview, 1-5

specifications, 4-9

ST registers, 2-3

Standalone disk testing, 1-7

Standalone SSD-E testing, 1-8

Standalone SSD-E/32i testing, 1-8

Stations, 6-8

Status

field bit assignments, 2-36

registers, 2-38

registers transfer instructions, 2-62

Storage, I/O buffer, 3-3

See also Disk drives; Disk subsystems; SSD-E

Subroutine libraries, 6-6

SUPERLINK/MVS product, 6-8
 Swapping. *See* Exchange: sequence
 Syntax
 CAL, 2-54
 with integer arithmetic instructions, 2-65
 System components, 1-1

T

T registers. *See* Intermediate registers
 Tape controller channel adapter, 3-7
 Tape drives and controllers, 1-6
 TCA-1 channel adapter, 1-6, 3-7
 TCA-2 channel adapter, 3-7
 TCP/IP, 6-7
 Telebit NetBlazer, remote support, 1-7
 Temperature
 monitoring. *See* WACS
 recommended water-supply, 1-17
 Test and set instruction, 2-3
 Training publications, 6-11
 Truncation. *See* Floating-point algorithms:
 multiplication

U

Unconditional branch instructions, 2-76
 Underflow condition. *See* Range errors,
 floating-point
 UniChem environment, 6-8
 UNICOS, 2-47, 6-1, 6-2
 UNIX, 6-1, 6-2
 UNIX station, 6-8
 Upgrades, with OWS-E, 1-8
 USCP protocol, 6-7
 UTC-1 channel adapter, 3-8
 Utilities
 general, 6-6
 UNICOS, 6-2

V

V register
 functions, 2-43–2-44
 general, 2-4, 2-7

 transfer instructions, 2-58, 2-61
 Values, special register, 2-53
 Vector
 chaining example, 2-45
 data, 2-4
 defined, 2-42
 examples, 2-43
 floating-point data, 2-4
 instructions, 2-44
 leading zero count instruction, 2-75
 length (VL) register, 2-7
 mask (VM) register, 2-7
 mask bits format, 2-49
 mask instructions, 2-10, 2-72
 memory references, 2-6
 population count instruction, 2-74
 processing, 1-1, 1-4, 2-39, 2-42–2-44
 segmentation and pipelining example, 2-41
 Vector functional units, 2-9–2-11, 2-40
 See also Floating-point: functional units
 Vector length (VL) register field, 2-37
 Vector length register transfer instructions, 2-62
 Vector mask (VM) register transfer instructions,
 2-62
 Vector operations, automatic, 2-43
 See also Floating-point: functional units
 VHISP channels, 1-5, 4-3
 VM station, 6-8
 VMEbus, 5-15
 Voltage monitoring. *See* WACS
 VT applications, 6-7

W

WACS, 1-7
 CRAY C916, 1-18
 CRAY C92A and CRAY C94A, 1-11
 CRAY C94 and CRAY C98, 1-14
 Warning and control system. *See* WACS
 WINS, 1-5
 Workstation interfaces. *See* WINS
 Write instructions, 2-64

X

XA register field, 2-37

Y

Y-MP mode. *See* C90 mode and Y-mode differences

Reader Comment Form

**Title: CRAY® C90 Series Functional
Description Manual**

Number: HR-04028-0A

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this manual?

- Troubleshooting
- Tutorial or introduction
- Reference information
- Classroom use
- Other - please explain _____

Using a scale from 1 (poor) to 10 (excellent), please rate this manual on the following criteria and explain your ratings:

- Accuracy _____
- Organization _____
- Readability _____
- Physical qualities (binding, printing, page layout) _____
- Amount of diagrams and photos _____
- Quality of diagrams and photos _____

Completeness (Check one)

- Too much information _____
- Too little information _____
- Just the right amount of information

Your comments help Hardware Publications and Training improve the quality and usefulness of your publications. Please use the space provided below to share your comments with us. When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

NAME _____
JOB TITLE _____
FIRM _____
ADDRESS _____
CITY _____ STATE _____ ZIP _____
DATE _____

[or attach your business card]



CUT ALONG THIS LINE

Fold



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



BUSINESS REPLY CARD

FIRST CLASS PERMIT NO 6184 ST. PAUL, MN

POSTAGE WILL BE PAID BY ADDRESSEE



**Attn: Hardware Publications and Training
890 Industrial Boulevard
Chippewa Falls, WI 54729**

Fold

STAPLE

**Cray Research, Inc.
Hardware Publications and Training
890 Industrial Boulevard
Chippewa Falls, WI 54729**