

CRAY-3 Hardware Description Manual

Publication Number: 3207

Fifth Revision

July 30th, 1993



**Cray Computer Corporation
1110 Bayfield Drive
Colorado Springs, CO 80906**

Copyright

Copyright © 1993 by Cray Computer Corporation. All Rights Reserved. This manual or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Cray Computer Corporation.

Copyright © 1986, 1987 by Cray Research, Inc. All Rights Reserved. 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, CFT, CFT2, CFT77, CRAY-2, CRAY X-MP, CRAY Y-MP2E, and SEGLDR are trademarks and CRAY, CRAY-1, CRAY Y-MP, HSX, UNICOS, and X-MP EA are registered trademarks of Cray Research, Inc.

bdb, CRAY-3, CRAY-4, CSOS, Doyle, Holmes, Hudson, stb, Watson, and Wiggins are trademarks of Cray Computer Corporation.

BSD is a registered trademark of the University of California, Berkeley.

Ethernet is a registered trademark of the Xerox Corporation.

Fluorinert is a registered trademark of 3M (Minnesota Mining and Manufacturing).

HYPERchannel is a trademark, and NSC is a registered trademark of Network Systems Corporation.

libtcl is authored by Professor John Osterhout, U.C. Berkeley and extended by Cray Computer Corporation.

Macintosh is a registered trademark and Quadra is a trademark of Apple Computer, Inc.

NeWS, NFS, OpenWindows, Sun, Sun Microsystems, Inc., SunView, and XView are trademarks and Sun Workstation and SunOS are registered trademarks of Sun Microsystems, Inc.

System V is a trademark, and OPEN LOOK and UNIX are registered trademarks of USL (UNIX System Laboratories) in the United States and other countries.

OSF and OSF/Motif are trademarks of Open Software Foundation.

POSIX is a trademark of The Institute of Electrical and Electronics Engineers, Inc.

SPARCstation and SPARCware are trademarks of SPARC International, Inc.

UltraNet is a registered trademark of Ultra Network Technologies, Inc.

VAST is a registered trademark of Pacific Sierra Research Corporation.

X Window System is a trademark of the Massachusetts Institute of Technology.

The CSOS operating system is derived from Cray Research, Incorporated's UNICOS operating system. The UNICOS operating system is derived from the USL's 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.

Reader Comments

If you have any comment about the technical accuracy, content, or organization of this manual, please tell us. You can contact us in any of the following ways:

Call our Technical Publications department at (719) 579-6464 from 8 a.m. to 5 p.m. (Mountain Time).

Send us electronic mail from a CSOS or UNIX system to pubs@craycos.com

Write to us at the following address:

Cray Computer Corp.
Technical Publications Department
P.O. Box 17500
Colorado Springs, CO 80935

We value your comments and will respond to them promptly.

Table of Contents

Copyright	i
Reader Comments	iii
Table of Contents	v
Introduction	ix

Chapter 1	CRAY-3 Overview	1
------------------	------------------------	----------

Packaging Overview	1	1.1
The Octants	4	1.1.1
The Modules	7	1.1.2
The Printed Circuit Boards	11	1.1.3
The Logic, Power and Resistor Plates	12	1.1.4
The Integrated Circuits	13	1.1.5
Design Overview	15	1.2
Performance Specifications	15	1.2.1
Common Memory	17	1.2.2
Background Processing	19	1.2.3
Foreground Processing	23	1.2.4
CRAY-3 Product Description	26	1.3

Chapter 2	Circuit Component Design	29
	GaAs Component Design.....	2.1
	Transistor Design.....	2.2
	Current Source Design.....	2.3
	Diode Design.....	2.4
	Capacitor Design.....	2.5
	Resistor Design.....	2.6
	Logic Package Design.....	2.7
	D Cell Design.....	2.7.1
	E Cell Design.....	2.7.2
	F Cell Design.....	2.7.3
	FF and FG Cell Design.....	2.7.4
	FJ and FK Cell Design.....	2.7.5
	Clock Amplifier Design.....	2.8
	Other Circuits.....	2.9
	The Vector Register Package.....	2.10
	The Silicon Memory Packages.....	2.11
	The CRAY-3 Cell Library.....	2.12
Chapter 3	Integrated Circuit Artwork	79
	The IC Assembly.....	3.1
	Routing an IC.....	3.2
	Verification of the IC Artwork.....	3.3
Chapter 4	Integrated Circuit Mask Production	101
	The Glass Plates or Masks.....	4.1
	Mask Data.....	4.2
	Fracturing.....	4.3
	Mask Layers.....	4.4
	Positioning of ICs from Mask to Wafer.....	4.5

Chapter 5	Integrated Circuit Fabrication	113	
	Integrated Circuit Fabrication	113	5.1
	Canon Wafer Stepper	116	5.1.1
	GaAs Fab Processing Steps	118	5.1.2
	Wafer Processing	124	5.2
	FAST Testing	124	5.2.1
	Grinding	125	5.2.2
	Dicing	126	5.2.3
	Pick and Place	127	5.2.4
	Die Bonding	128	5.2.5
Chapter 6	Module Structure	133	
	Overview.	133	6.1
	Printed Circuit Board Manufacturing.	139	6.2
	The 25 mm Circuit Boards	141	6.3
	Logic Plates	150	6.4
	Resistor Plates	154	6.5
	Power Plates	158	6.6
	Memory Board Bypass Capacitors.	160	6.7
	Power Blades	162	6.8
	Logic Connectors	167	6.9
Chapter 7	Module Assembly, Testing and Repair	171	
	Final Assembly of the Module	171	7.1
	Kitting	173	7.1.1
	Board Soldering	174	7.1.2
	Die Attach	177	7.1.3
	Twist Pin Manufacturing	180	7.1.4
	Stacking the Module Assembly	183	7.1.5
	Pinning the Module Assembly	186	7.1.6
	Module Testing	190	7.2
	Pulse-Power Testing	190	7.2.1
	Full-Power Testing	196	7.2.2
	Module Repair	202	7.3

	Failure Analysis	203	7.4
Chapter 8	Cabinet Design	207	
	System Cabinet	207	8.1
	Octant Assembly	218	8.2
	Power Supplies	222	8.3
	Wire Harnesses	228	8.4
	Cooling	231	8.5
	Translator Cards	235	8.6
	System Control Pod	237	8.7
	Front Panel User Controls	241	8.7.1
	Maintenance Panel	244	8.7.2
	Electronic Module	245	8.7.3
	Motor Starters	247	8.7.4
	Clock Source	247	8.7.5
	AC Box	247	8.7.6
	DC Bulkhead	248	8.7.7
	Peripheral Equipment	251	8.8
	DD-49 Disk Unit	251	8.8.1
	DS-40 Disk Subsystem	251	8.8.2
	CRAY-3 Circuits for the DD-49 and DS-40	252	8.8.3
	RAID Disk Array Subsystem	257	8.8.4
	CRAY-3 HIPPI Circuitry	258	8.8.5
Appendix A	Integrated Circuit Packages Used in the CRAY-3	259	
Appendix B	CRAY-3 Board Stack Assignments	279	
	List of Figures	293	
	Index	299	

Introduction

This is the fifth revision of this manual. The previous version was dated January 15, 1991. All sections have been revised and many new sections added to reflect changes which have occurred over the last twenty-nine months in the design, manufacturing and testing of the CRAY-3 supercomputer system. This current version of the manual has over 300 pages, more than 160 illustrations and 22 tables, whereas the previous version had just over 100 pages and 45 illustrations.

Much effort has been expended to make this latest edition of the manual accurate and up-to-date. However, this is really an impossible task since procedures and design are in an almost continuous state of flux. At some point a decision must be made to “go to press” knowing that by the time the manual reaches most readers certain items of fact will have already changed.

Purpose

The CRAY-3 Hardware Description Manual is intended to provide the person who is working on the project with a general reference manual to the component hardware that comprises a CRAY-3 computer system. It is ideal for introducing a new employee to the overall project.

The manual will also find interested readers among current and future customers, since it provides a unique summary of information on the CRAY-3 that is not available in any other form.

Organization

The manual first presents an overview of the CRAY-3 computer system looking at both the overall architecture of the machine and the implementation of that design in the hardware and packaging. This is done in Chapter 1. The overview of the packaging begins at the highest level (the System Cabinet) and works down to the lowest level (the integrated circuit die). This helps the reader become familiar with how all the parts fit together in the complete machine. The overview of the packaging is presented before the overview of the design to introduce some of the terminology which is employed in discussing the design.

In the remaining chapters each major component is discussed in much greater detail. In these chapters the presentation order is just the opposite of that in the introduction. The chapters proceed from the lowest level (the design of the integrated circuit components) to the highest level (the system cabinets and peripheral equipment).

The manual has tried to reach a balance between being overly technical and obtuse, and simplistic and obvious. This has not been an easy task given the inherently technical subject matter and the broad spectrum of potential readers.

Hopefully, employees of Cray Computer Corporation, whatever their position, will find the CRAY-3 Hardware Description Manual interesting and informative, for they are viewed as the primary readers.

CRAY-3 Overview

This chapter introduces the CRAY-3 computer system by providing three overviews of the machine. The first overview looks at the implementation of the design in the hardware and packaging. Giving an overview of the hardware first will introduce some of the terminology needed for the following discussions. The second overview looks at the design and architecture of the machine. This overview includes a summary of the machine's performance specifications. Much of this discussion uses the CRAY-2 as a point of reference for performance comparisons. The third overview ends the chapter with a summary of the machine in the form of a product description.


It is important for the reader to grasp the big picture presented in these three overviews so that later discussions about each individual component or group of components can be better understood in the overall context of the complete machine.

1.1 Packaging Overview

The CRAY-3 hardware is constructed of synchronous networks of binary circuits. These circuits are packaged in 336 modules for a 16-processor machine. The 336 modules fit into the top section of the System Cabinet. The System Cabinet for a 16-processor machine is in the shape of an octagon 109.22 cm (43 inches) wide and 121.92 cm (48 inches) above a computer room false floor. The power supplies for the modules are located below the modules

in the bottom section of the System Cabinet. Another cabinet, the Control Pod (commonly called the C-Pod), located near the System Cabinet, contains the pumps, heat exchangers, reservoir, filters and control circuits for the cooling system, the system clock, controls for the power supplies and other mechanical and electronic monitoring devices. Computer control and program initiation and monitoring is done from a workstation console which may be located in another room.

FIGURE 1. Four-Processor System Cabinet and Control Pod.




The System Cabinet and the C-Pod are located near each other in a computer room with a false floor. The cooling and electrical connections required between the two cabinets are made underneath the false floor in the computer room so as to be invisible. The System Cabinet (the foreground cabinet in Figure 1) contains all of the logic and memory circuits of the computer in the top eight inches. The remainder of the System Cabinet houses all of the power supplies for the logic and memory circuits of the machine.

The CRAY-3 represents the first supercomputer to use gallium arsenide integrated circuits instead of silicon integrated circuits for all of its logic

circuits. Silicon integrated circuits are still used for the memory circuits in the machine.

The Four-Processor CRAY-3 System delivered to the National Center for Atmospheric Research on May 24, 1993 uses 62,738 integrated circuits all located on the 84 modules directly underneath the bronzed acrylic lid. The machine requires 90,000 watts of power and gives off 310,000 BTUs of heat—enough to heat six 2,000 square foot homes.

FIGURE 2. Eight-Processor System Cabinet.



All of the binary circuits in the machine are synchronized to a 500 megahertz oscillator located in the C-Pod. The oscillator signal is transmitted as a square wave over 60 ohm twisted-pair wires to each of the module connectors in the System Cabinet. Wire lengths are controlled so that the travel time to the individual modules is accurate within 50 picoseconds. The oscillator square wave is delivered to each individual circuit package requiring a clock input within each module through traces on the logic plates and printed circuit

boards. A pulse is then formed from the square wave by the clock amplifiers in each integrated circuit package to gate data and control into the latches within the IC packages. This strobe pulse occurs simultaneously throughout the machine with a period of two nanoseconds. This time is referred to as the "machine clock period".

Another pulse is also formed from the inverted clock signal by identical clock amplifiers and is used to gate information into the latches one-half clock period later than the normal clock signal. This half clock period of time (one nanosecond) is referred to as a "clock phase time". The normal clock signal latches information in the even clock phases. The invert clock signal latches information in the odd clock phases. This dual phase system clock allows information to be latched into successive logic cells every nanosecond.


The System Cabinet is covered by a cosmetic shell or skin which gives the CRAY-3 the appearance of being of one piece. However, the System Cabinet is really modular in design and consists of groups of octants. Each octant contains a one- or two-processor piece of a complete 16-processor machine, including the memory modules, I/O module or modules and power supplies associated with the modules for that octant. Each module is in turn made up of various types of printed circuit boards; namely, the power plates, the resistor plate, the logic plates, and the smaller printed circuit boards which hold the IC packages for both logic and memory circuits.

1.1.1 The Octants

The 336 modules for a 16-processor machine are arranged in the System Cabinet in eight columns which form a closed octagon. Each column or octant contains 42 modules in two ranks, arranged in an outer rank of 32 and an inner rank of 10. This arrangement of modules is illustrated in Figure 3. Fluorinert, an inert electronic liquid, circulates in the cabinet frame and flows through the modules for cooling. The temperature, liquid velocity and turbulence through the modules are all controlled by various sensors and monitoring devices. The octagon of module columns is located on top of a similar structure containing power supplies for the system. The power supplies are also cooled by the circulating liquid. Total power consumption for the system is approximately 360 kilowatts or approximately 45 kilowatts per octant.

The CRAY-3 computer system can be configured in one-, two-, four-, eight-, or 16-processor versions. Since each one- or two-processor segment of the machine resides in one of the octants, a one- or two-processor machine would consist of one octant, a four-processor machine would consist of two octants, an eight-processor machine would be made up of four octants, and a 16-processor machine would have a full complement of eight octants utilizing the full octagon for its modules and power supplies.


FIGURE 3. Module Placement In the CRAY-3.



Before leaving our introduction to octants, one should note that the term "octant" is often used to refer to two different, but integrally related parts of the CRAY-3. The term can be used to describe that physical portion of the System Cabinet where a group of 42 modules and their associated power supplies and wire harnesses reside. This is the sense in which we have used the term above. In this sense the term is often expanded to "octant assembly" or shortened to simply the "tank." This terminology is most often employed by the mechanical design group. Alternately, the term "octant" is often used by the logic design engineers when referring to a specific group of modules which could be placed into an octant assembly. In this sense one may refer to "Octant 6" (meaning a specific group of 42 modules) or to "an octant of memory" (meaning a specific

group of 32 memory modules) which could move about from one System Cabinet's octant assembly to another, or from one two-processor pulse power test station to a System Cabinet (or "tank"). In the first sense, a specific, physical location is being intended. In the second sense a specific group of modules that form a functional unit and that may move from one specific location to another is being intended.

FIGURE 4. Module Types Used in a CRAY-3 Octant.




1.1.2 The Modules

There are four basic types and ten unique modules used in the CRAY-3, each designated by a letter of the alphabet:

- The Background Processors—A, B, C, and D Modules.
- The Foreground Processor—E Module.
- Common Memory—F, G, and H Modules.
- I/O Modules—L and M Modules.

The Background Processors (A, B, C and D Modules) make up one processing unit (similar to a CPU in smaller computers). When we say “this is a two-processor CRAY-3” we are implying that it has two A, two B, two C, and two D modules. Likewise, a 16-processor machine has sixteen A, sixteen B, sixteen C and sixteen D modules. Each machine has one E module (the Foreground Processor) and each octant will also have 32 Common Memory modules, and zero, one, or two I/O modules of some kind depending upon customer needs. Therefore, in a 16-processor machine with 336 modules, 256 are reserved for Common Memory, 64 for Background Processors, one for the Foreground Processor and 15 for control of I/O devices such as disk drives.


FIGURE 5. CRAY-3 Electronic Module.



The modules are three dimensional structures measuring 121 mm by 107 mm by 7.14 mm thick. The circuit packages within the modules are unpackaged gallium arsenide or silicon die mounted directly on multi-layer circuit boards.

These boards are arranged in a four-by-four matrix of board stacks, each stack four boards deep (thick). In the center of this structure are two multi-layer logic plates which provide routing of logic signals between the sixteen board stacks in the module. Between the two logic plates of the module are two multi-layer power plates for power distribution to the board stacks. At the center of the module, between the power plates, is the resistor plate for signal termination. Some experiments are currently being conducted which place the resistor plate on the outside of this inner group of plates to provide for better cooling of the resistors. Other experiments, which also have as their goal better cooling of the resistors, use individual resistor boards on the outside of the logic plates rather than one resistor plate. The layers of the module sandwich can be seen in Figures 6 and 7.


FIGURE 6. Layers in a Module.



1.1.3 The Printed Circuit Boards

The printed circuit boards used in the CRAY-3 have the same basic properties of most modern printed circuit boards. They have very thin copper traces and serve as a place for mounting all of the ICs. However, they are smaller and thinner than most printed circuit boards commonly seen in electronic equipment. CRAY-3 printed circuit boards are only 28 mm x 25 mm in size, have eight layers (two of which are used for logic signal trace lines) and are only 0.38 mm thick.

FIGURE 9. Sizes of the Basic Components in a CRAY-3 Module.



Illustrations shown are approximately actual size.

Up to sixteen GaAs die may be attached to each of these small circuit boards with fifty-two very small gold pins used to attach each die. The printed circuit boards provide the copper traces needed to interconnect the signals from one die to another, as well as to distribute the power and logic signals from the

power and logic plates to the die. The logic signal trace lines on each printed circuit board are only 0.048 mm wide (a human hair averages 0.070 mm in diameter). The digital data for the 25 mm circuit boards is routed, drawn, and tested in the Artwork Department at Cray Computer Corporation. The circuit boards are manufactured at Cray Computer Corporation's own facility in Colorado Springs.

1.1.4 The Logic, Power and Resistor Plates

The Logic Plate is simply a larger size circuit board (107 mm by 104 mm and 0.46 mm thick). The reason for this difference in size should be apparent if you think of the logic plate's primary function—to distribute logic signals to and from the smaller 25 mm circuit boards which fit side-by-side making a square four circuit boards by four circuit boards, or sixteen in each layer of circuit boards. So the logic plate needs to cover the area of four 25 mm circuit boards if it is going to distribute signals to all those smaller boards. The logic plate contains eight layers, as do the 25 mm circuit boards, but whereas only two layers are used for drawing logic trace lines on the circuit boards, six layers are used for drawing logic trace lines on the logic plates.

When people refer to "the logic plate" they are often using that expression to describe more than a single logic plate since the logic plates are always used in pairs, drawn together, and considered a single, functional unit even though they are physically two separate plates. Furthermore, in addition to the two logic plates, the two power plates and the single resistor plate are also generally considered to be part of the logic plate's working assembly. Therefore, in our discussions, the power plates and resistor plate will be included in the section on the logic plates.

The power plates contain no logic signal traces. The power plate serves to deliver the necessary voltages from the power blades to all of the power pins at the corners of each 25 mm circuit board. From there the power layers of each circuit board deliver the appropriate voltages to the proper pins of each IC. The power plate is longer than the logic plates in one dimension to provide space for the power blades to contact the outside surface of each power plate.

The resistor plate is sandwiched between the two power plates and contains the 55 ohm terminating resistors used to match the impedance of the signal transmission lines used throughout the logic circuits of the computer. The resistors are connected to the appropriate signal trace line through twist pin jumpers which carry the signal vertically through the module layers. As noted earlier, experiments are under way which change the positioning of the resistor plate in the module sandwich or which replace the resistor plate with individual resistor boards, also changing the position of the resistors relative to

other layers of the module sandwich. The outcome of these experiments will be reflected in later revisions of this manual.

1.1.5 The Integrated Circuits


The smallest electronic parts used in the CRAY-3 are the integrated circuit packages or ICs. Sometimes people will refer to them as “chips” or generally after they have been manufactured and without being packaged in tiny containers, as “die.” The physical size of the gallium arsenide die as they appear on the printed circuit boards is 3.835 x 3.835 millimeters. There are 498 different GaAs integrated circuit packages used in the logical networks of the machine. Each of these circuits consists of an array of up to 128 basic logic cells. There are 36 different logic cells from which all the circuits are built. The maximum equivalent gate capacity of the circuit packages is approximately 500. The logic circuit packages have 52 bonding pads for connection to the printed circuit boards.

The ICs are the active circuits that actually perform the logical functions which enable the computer to perform its calculations. It should not be surprising then that there are a lot of ICs in a CRAY-3 computer system. The 16-processor machine contains over 266,000 ICs of which 147,456 are used for SD Common Memory. A two-processor machine has over 29,000 integrated circuit packages. Since there are only 498 unique types of ICs, many are used again and again throughout the machine to perform the necessary logical functions required. For example, the Vector Register integrated circuit package is used over 3,000 times in a 16-processor CRAY-3.

Nearly all other parts of the CRAY-3 are there simply to serve these tiny devices—to give them a safe place in which to reside, to keep them cool, to give them the right kind of power and to provide communication lines so they can talk to each other and communicate with the outside world.

The diagram in Figure 10 illustrates how the basic components of the CRAY-3, from the ICs to the Modules, fit together to make up a complete machine.

FIGURE 10. Basic Components of the CRAY-3.



1.2 Design Overview

The CRAY-3 computer system represents a major enhancement over the CRAY-2 computer system, providing an order of magnitude performance improvement at a comparable cost. This performance gain is achieved through the use of three new design features 1) the use of gallium arsenide logic circuits in place of the silicon circuits previously used, 2) the use of innovative packaging and cooling technologies and 3) the implementation of a number of architectural changes, including more processors and enhanced memory access.

1.2.1 Performance Specifications

Each processor in the CRAY-3 system is slightly more than two times as fast as its CRAY-2 counterpart. With a four-fold increase in the number of total processors, this leads to an overall improvement range better than eight times that of a CRAY-2.

The CRAY-3 system is functionally equivalent to the CRAY-2 system. Logic enhancements to the system have enabled performance to be increased while providing an extension of the original CRAY-2 instruction set. The following table lists the major differences between the CRAY-2 and CRAY-3 systems:


TABLE 1. Performance Comparisons Between the CRAY-3 and the CRAY-2.

Parameter	CRAY-3	CRAY-2
Circuit Speed	2 Nanosecond Clock (500 MHz) Eight Levels of Logic per Clock	4 Nanosecond Clock (250 MHz) Four Levels of Logic per Clock
Processors	Up to 16 Background Processors One Foreground Processor	Four Background Processors One Foreground Processor
Instruction Issue Rate	Each Clock Period Maximum Instruction Issue Rate of 8000 MIPS (Million Instructions Per Second) for a 16 Processor Configuration.	Every Other Clock Period Maximum Instruction Issue Rate of 500 MIPS for a 4 Processor Configuration.
Common Memory	512 Million Words of Static Semiconductor Memory or 1024 Million Words of SRAM or 2048 Million Words of SRAM (with memory enhancement now under development).	256 Million Words of Dynamic Semiconductor Memory.
Memory Chip Access Time	22 Nanoseconds	80 Nanoseconds
Common Memory Bandwidth	16 Gigawords Per Second	1 Gigaword Per Second
Common Memory Access Ports	Each Background Processor has a Bi-Directional Access Port. Reads and Writes may Proceed Simultaneously.	Each Background Processor has One Access Port. Read or Write but not both Simultaneously.
Floating Point Computation	Maximum Burst Rate of 16 Gigaflops	Maximum Burst Rate of 2 Gigaflops
Vector Register Structure	Eight Vector Registers per Background Processor. Employs Tailgating: writing may begin immediately after previous data reading is initiated.	Only two late versions of the CRAY-2 employed Tailgating. Both versions were transferred to Cray Computer Corporation at inception.

1.2.2 Common Memory

The CRAY-3 system Common Memory consists of 512 storage banks of up to four million words each. Each word consists of 64 data bits and eight error correction bits. This memory is shared by the Foreground Processor, Background Processors and peripheral equipment controllers. It contains program code for the Background Processors as well as data for problem solution. System tables are located here for the Foreground Processor but foreground program code is not. Data buffers for the disk files are also located in Common Memory.

FIGURE 11. Common Memory Stack Usage.



There are two memory banks on each Common Memory module. Each memory bank occupies six out of the sixteen stacks on a module. The remaining four stacks are control stacks (see Figure 11). A memory bank can utilize 16 background processor memory access ports. Total memory bandwidth is 128 gigabytes per second. Total memory capacity is eight

gigabytes. Each Background Processor can read and write a word of data per clock period (two nanoseconds) in a vector mode.

The silicon Static Random Access Memory integrated circuits used in Common Memory each contain either 262,144 or 4,194,304 bits of data (with enhanced memory). A memory bank consists of 288 of the 256K-by-1 SRAM integrated circuits or 72 of the 4M-by-1 SRAM integrated circuits. A memory module then contains 576 x 256K or 72 x 4M bits of data plus logic circuits to support the memory access paths. Memory access time at the circuit level is 22 nanoseconds. Memory cycle time is 30 nanoseconds.

The memory enhancement for the CRAY-3, under development as this revision was published, would increase the Common Memory of the machine four times per octant of memory. With the 4 meg SRAM circuit enhancement a four-processor CRAY-3 (two octants) would contain 512 million words of Static Random Access Memory. A 16-processor machine would contain two gigawords of Common Memory.

The CRAY-3 GaAs integrated circuits are significantly faster than the memory circuits. This speed discrepancy is used to minimize the number of physical data paths that are necessary to connect the 512 memory bank modules to the 32 memory ports. Data is transferred by utilizing an 18-bit packet. The 64-bit data word and eight error correction bits require a four clock period transmission time—18 bits per clock period. There are independent read and write packets between the memory bank and access port. Memory bank address utilizes a 12 bit packet. The 24 bit internal bank address requires a two clock period transmission time between the access port and memory bank—12 bits per clock period.

An eight bit error correction code is generated at the memory access port as the write data is transmitted to the memory bank. This code is interpreted at a readout port for Single Error Correction and Double Error Detection, often abbreviated as SECDED. This allows the computer to continue to properly transfer data to and from common memory even if several single bit errors occur in the data during the transfer or storage process. Double bit errors are also detected but cannot be corrected.

1.2.3 Background Processing

The Background Processors in the CRAY-3 system are composed of processing elements similar to those in a CRAY-2 processor. The CRAY-3 instruction set includes similar registers and arithmetic functions, retaining all of the CRAY-2 instructions. Some of the instruction source and destination register designators have been changed to provide consistency and logical simplification.

Eight new instructions have been added as summarized in the following table:


TABLE 2. New Instruction Codes for the CRAY-3.

Machine Code	CAL Instruction		Description
001---	CMR		Complete memory references.
023ij1	A_i	A_j+1	Enter A_i , increment A_j .
035-j-	RT	S_j	Copy S_j to RT (Monitor Mode only).
074i--	V_i	[lm]	Direct read V_i from Local Memory.
075i--	[lm]	V_i	Direct write V_i to Local Memory.
134ijk	V_i	$\langle A_j, A_k \rangle$	Two-port read V_i from CM (A_j) stride A_k .
135ijk	$\langle A_j, A_k \rangle$	V_i	Two-port write V_i to CM (A_j) stride A_k .
177---	PASS		No-op.

Computation in a background mode implies a memory-to-memory computation. A job is initiated by the Foreground Processor. Data is moved from the disk files to the semiconductor memory under Foreground Processor supervision. The program code for the Background Processors is positioned in the Common Memory and the computation field is defined. The Foreground Processor then initiates the background computation using one or more Background Processors as required. The Background Processors may call for further peripheral activity through the Foreground Processor as the computation proceeds.

Figure 12 contains a diagram illustrating the architecture of a CRAY-3 Background Processor.

FIGURE 12. Block Diagram of a Background Processor.



Each background processor has a small high speed local memory to hold scalar or vector operands during a computation (16,384 words are available for each Background Processor). Data is moved from the Common Memory to the local memory and returned at the end of each computation. Arrays of data are addressed by the Background Processors directly in the Common Memory. The access to data common to multiple Background Processors is interlocked by the Background Processor semaphore flags.

Resources of a Background Processor are summarized in the following table:

TABLE 3. Resources of a CRAY-3 Background Processor.

Functional Division	Function Performed	Specification
Eight S Registers	Scalar operands	64 bit length
Eight A Registers	Address and integer operands	32 bit length
Eight V Registers	Vector segments	64 elements of 64 bits
Instruction Buffer	Instruction loops	512 parcels of 16 bit length
Local Memory	Address, Scalar and Vector register backup	16,384 words of 64 bit length
Floating Point Functional Units	Computation in 64 bit floating point mode	
Integer Functional Units	Computation in 64 bit integer mode	
Address Functional Units	Computation in 32 bit integer mode	
Vector and Scalar Logical Units	64 bit logical computations	
Vector and Scalar Shift Units	Shifts, population, and leading zero counts	

Figure 13 shows how the different functions of each Background Processor are distributed among the four modules which comprise each Background Processor.

FIGURE 13. Background Processor Functional Divisions Among the Modules.


ONE BACKGROUND PROCESSOR			
A Module	B Module	C Module	D Module
Scalar Registers Bits 00-15	Scalar Registers Bits 16-31	Scalar Registers Bits 32-47	Scalar Registers Bits 48-63
Vector Registers Bits 00-15	Vector Registers Bits 16-31	Vector Registers Bits 32-47	Vector Registers Bits 48-63
Local Memory Bits 00-15	Local Memory Bits 16-31	Local Memory Bits 32-47	Local Memory Bits 48-63
Integer Add, Logical, Mask, RTC Bits 00-15	Integer Add, Logical, Mask, RTC Bits 16-31	Integer Add, Logical, Mask, RTC Bits 32-47	Integer Add, Logical, Mask, RTC Bits 48-63
Address Registers	Branch Control	Floating Multiply	Instruction Issue Control
Address Multiply	Program Register	Reciprocal Approximation	Instruction Buffers
Address Add Scalar Shift	Floating Add		Foreground Interface
Vector Length Register			Vector Shift
Common Memory Address	Common Memory Write		Common Memory Read

1.2.4 Foreground Processing

The Foreground Processor supervises overall system activity and responds to requests for interaction between the system members. System communication is accomplished through four, high-speed synchronous data channels. These channels interconnect the Background Processors, Foreground Processor, disk control units and host system interfaces. The Foreground Processor does not execute the system program code.

The Foreground Processor code is loaded at deadstart from a file on the maintenance console. This memory then becomes a read only memory. It cannot be altered during the operation of the system. Data for supervision of the system is maintained in the Common Memory and is moved to the Foreground Processor's local data memory as required.

FIGURE 14. Data Channels (One of Four).




The primary function of the Foreground Processor program is to poll the four communication channel loops for requests from the Background Processors and control the communications on the four channel loops to facilitate data transfers between controllers on the channels. Many of the requests that will be recognized by the Foreground Processor require responses within a microsecond or less.

Since the majority of Foreground Processor activity involves data transfer between the disk file storage units and the common semiconductor memory, the system provides for a mixture of 12 megabytes per second disk file storage

unit interfaces and 100 megabytes per second interfaces to main storage sub-systems such as disk arrays.

FIGURE 15. Block Diagram of the Foreground Processor.



The Foreground Processor is divided into seven operational sections: 1) deadstart circuits, 2) instruction issue mechanism, 3) local data memory, 4) channel communication, 5) functional units, 6) console communications and

7) real time clock. Resources of the Foreground Processor are summarized in the table below:

TABLE 4. Resources of the Foreground Processor.

Functional Division	Functional Purpose	Specification
Local Data Memory	Temporary storage	4096 words of 32 bit length.
Instruction Memory		65,536 bytes of 8 bit length arranged in 16 banks.
Integer Functional Units	Computation in 32 bit integer mode.	
Four Communication Channels		One gigabyte per second each.

1.3 CRAY-3 Product Description

We will conclude Chapter 1 by giving a brief product description of the CRAY-3 supercomputer system. This description will summarize the basic features of the CRAY-3 in terms of hardware, software, processor capabilities and peripheral and network interfaces. This summary of the CRAY-3 is given in Table 5, below:

TABLE 5. CRAY-3 Features Summarized.

Hardware	
Feature	Specification
Background Processors (CPUs)	1, 2, 4, 8 or 16
Memory Size	128, 256, 512, 1024 or 2048 Million Words 1 Bank per Million Words up to 512 64 Bit Words with SECDED
Foreground Processor	1 (I/O and System Control Processor)
External I/O Channels	8 to 60
I/O Transfer Rate	Up to 4 Gigabytes per Second
Clock	2 Nanosecond (500 Megahertz)
Memory Transfer Rate	1 Gigaword per Processor per Second 16 Gigawords per second for a full system
Peak Performance	1 Gigaflap per Processor 16 Gigaflaps for a full system

Software	
Feature	Specification
Operating System	Extended UNIX
Compilers	Fortran 77 Standard C Both have extensions for Vectors and Multi-tasking.
Networks	TCP/IP
User Tools	Editors Debugger

Background Processor	
Feature	Specification
Functional Units	3 Vector 3 Scalar 2 Floating Point (Shared Vector/Scalar) 1 Address
Registers	8 Vector 1 Vector Mask 1 Vector Length 8 Scalar 8 Address 1 Program 1 Base 1 Limit
Instruction Stack	128 64-Bit Words
Real Time Clock	64 Bits
Local Memory	16,384 Words
Semaphore Flags	17
Ports to Memory	2 (1 Input and 1 Output)

Peripheral and Network Interfaces	
Feature	Specification
Networks	HIPPI (100 MBytes per Second) HYPERchannel by Network Systems Corp. VME via Cray Research, Inc. FEI-3 Ultra Technologies via HIPPI
Disk Storage	RAID Disks via HIPPI Cray Research, Inc. DD-49 and DD-40

To enable a machine to perform billions of operations every second you must have thousands upon thousands of switching circuits that are very fast and very close together. The design and packaging of the CRAY-3 is certainly a case of form following function with a logic gate density of 96,000 per cubic inch. In the next chapter we turn our attention to the smallest building blocks that make up the CRAY-3: the components inside the CRAY-3 GaAs integrated circuit packages.

Circuit Component Design

This chapter first presents a summary of some of the principles involved in designing the very small electronic components that make up the Gallium Arsenide integrated circuits used in the CRAY-3. It continues with a discussion of the basic logic cells used to construct a CRAY-3 integrated circuit. The chapter concludes with a discussion of the CRAY-3 Cell Library.

2.1 GaAs Component Design

Circuit component design in GaAs is a relatively open affair. In the silicon world, circuit design usually involves an array of gates in which the diffusion set is a constant over a large number of circuit types (gate array based). A metallization variation provides the unique combination of logic functions required. In our use of gallium arsenide, the design begins with the basic substrate for each unique circuit package (standard cell based).

The field effect transistors used in GaAs circuits are slightly different from the equivalent transistors in silicon circuits. The reason for this difference is that in silicon a metal oxide is formed over the substrate. In GaAs this is not the case since it is difficult to grow a stable oxide. Hence, field effect transistors in silicon are Metal Oxide Semiconductors (MOSFETs) while FETs in GaAs are Metal Semiconductors (MESFETs). The gate metal, which causes the field effect, is also a Schottky diode in GaAs circuits. This means that gate current will occur as the gate voltage becomes positive with respect to the source by

about 600 millivolts. In the case of silicon, the gate metal is not a diode and no current flows.

Circuit design for the CRAY-3 logic circuits involves eight levels of masks. There are actually a total of ten masks used in the GaAs foundry process. However, the first mask (alignment marks) and the last mask (passivation openings) are not unique to the individual circuits and are reused over the entire logic family. The fifth mask (proton isolation) is derived from the second and third mask and requires no additional design. The masks which are unique to the individual circuit design are assigned a letter designation for each mask level. These letter designations and their meaning are shown in the following table:

TABLE 6. Layers Used in the Design of CRAY-3 IC Components.

Mask Level	Process and Function	CrayDraw Layer	Standard CrayDraw Color
Level A	N- Implant. Semiconducting Channel.	Layer 1	Red
Level B	N+ Implant. Essentially terminates the semiconducting channel.	Layer 2	Green
Level C	Ohmic Metal. Metal terminals for the elements of a device.	Layer 3	Yellow
Level D	First Layer Metal. Used for connections between devices. Forms a Schottky Diode without Ohmic Metal.	Layer 4	Blue
Level E	Via Windows. Used to make openings in the dielectric between First Layer Metal and Second Layer Metal so the two layers can be connected electrically at that point.	Layer 5	Magenta
Level F	Second Layer Metal. Used for connections between devices.	Layer 6	Cyan
Level H	Recess Insulator. Makes an opening in the dielectric (insulator) so the gate metal can be lowered to the semiconducting channel.	Layer 8	Tan

Figures 16 through 20 illustrate the design of a 4.0 micron wide transistor with an individual level shaded in each of the five drawings. A 4.4 micron field effect transistor will represent the smallest FET used in the CRAY-3. All

devices are drawn 0.4 microns wider than needed to compensate for the shrinking of 0.4 microns during the fabrication process. All drawings of CRAY-3 ICs are made on a 0.1 micron grid using CrayDraw (an in-house custom designed Computer Aided Design program) or Cadence (a commercial CAD program). Only objects which can be made of rectangular structures with edges on the grid are allowed.

2.2 Transistor Design

A field effect transistor consists of a source element, a gate element and a drain element. Electrons flow through a semiconducting channel from the source to the drain. The gate controls the amount of current which flows through the semiconducting channel. The design of a 4.0 micron wide transistor is illustrated in Figures 16 through 20. The source and drain elements are symmetrical and the drain is determined by whichever element has the higher external voltage applied.

The semiconducting channel is illustrated in Figure 16. The width of the channel is the vertical dimension of the shaded area as drawn. This is by definition the width of the transistor. In this case the channel is drawn 4.4 microns wide. The length of the channel is the space from the N+ implant on the source side of the gate to the N+ implant on the drain side of the gate. All transistors used in CRAY-3 circuit designs have a channel length of 2.2 microns (see Figure 17).

FIGURE 16. 4.0 Micron Transistor—Semiconducting Channel.




Figure 17 has the areas of N+ implant shaded. The sheet resistance of the N+ implant is approximately 400 ohms per square and the sheet resistance of the N- implant is approximately 440 ohms per square. The double implanted area essentially terminates the effective transistor channel. The N- implant continues the length of the N+ implant and overlaps it. This makes the implant under the ohmic metal as highly doped as possible.

FIGURE 17. 4.0 Micron Transistor—N+ Implants.






Figure 18 shades the areas of the transistor design which have ohmic metal. The ohmic metal breaks down the substrate surface barrier which would otherwise result in a Schottky diode. These two areas of ohmic metal are essentially the metal terminals of the transistor for the source and the drain elements.

FIGURE 18. 4.0 Micron Transistor—Ohmic Metal.




Those areas of the transistor design which have first layer metal are illustrated in Figure 19. These metal paths on the source and drain connections of the transistor are 16 grid units wide, or 1.6 microns in width, and are used to connect the device to other portions of an integrated circuit package. The gate metal is narrowed to 0.8 microns as it passes over the N- implant channel to reduce the electrical capacitance of the gate metal to the channel. This capacitance is one of the components which determines the speed of the transistor.

FIGURE 19. 4.0 Micron Transistor—First Layer Metal.




A recess opening, which is shown in Figure 20, is made through the dielectric over the channel. The gate metal is lowered to the channel implant through this opening to make the device. All other metals except ohmic metal stay on top of the dielectric to reduce current leakage paths to the GaAs substrate.

FIGURE 20. 4.0 Micron Transistor—Recess Opening.



The electrical characteristics of a transistor are shown in the graph in Figure 21. The horizontal coordinate is the drain voltage with respect to the source voltage. The vertical coordinate is the amount of current that flows in the drain as a function of drain voltage and gate voltage.

FIGURE 21. Electrical Characteristics of a Transistor.



The graph shows five curves of drain current versus drain voltage for five constant gate voltages. The curve labeled $V_{gs} = 0$ volts has special significance because this is the curve applicable to transistors used as current sources. The point on this curve where the drain voltage is 2.5 volts has additional significance. The current at this point is chosen as a parametric reference for the amount of current capacity in the transistor. This current value has a special label (I_{dss}). I_{dss} stands for the drain current (I_d) with the gate voltage at the source supply voltage (ss).

The gate voltage which causes the drain current to become zero also has a special name. This is called the transistor pinch-off voltage. Gate voltages more negative than the pinch-off voltage have no additional effect. There is a


relationship between the pinch-off voltage as defined above and the drain voltage which causes the curves in the graph to break and become more or less flat. This is approximately the same voltage value. The design pinch-off voltage for the CRAY-3 transistors is -0.55 volt. Actual values range from -0.48 volt to -0.62 volt.

2.3 Current Source Design

Current sources can be used instead of resistors in many design applications. The current sources give a sharper transition to the electrical signals than a resistor and usually take less space. A current source is essentially a transistor with the gate element tied electrically to the source element. The electrical characteristic of the current source is therefore limited to one of the curves in the graph in Figure 21. This is the drain current versus the drain voltage curve with the gate voltage zero.

The design of a 4.0 micron wide current source is illustrated in Figure 22. Most of the description of transistor design applies to this device as well. The same five drawing types illustrated in Figures 16 through 20 would illustrate the five levels of process masks used in a current source design. The only difference between the current source design illustrated and the 4.0 micron transistor design previously described is the connection of the gate metal to the source metal. This is best shown when the first layer metal is shaded, as in Figure 22.

FIGURE 22. 4.0 Micron Current Source—First Layer Metal.




2.4 Diode Design

Diodes are easy to come by in gallium arsenide circuit design. Whenever first layer metal passes over an implanted area of the substrate a Schottky diode is formed. The electrical characteristics of the diode depend on the size of the metal overlapping the implant area and the resistance of the implant path to an ohmic contact.

There are two styles of N+ only diode configurations used in CRAY-3 circuit design. These are labeled type A diodes and type B diodes. The type A diode has a path resistance of 200 ohms per square micron of diode area. The type B diode has a path resistance of 650 ohms per square micron of diode area. The type A diode is used exclusively as the logic forming input device for cell structures (the logic gates). The type B diode is used for level shifting and clamping functions.


The design of a type A diode is illustrated in the five drawings in Figures 23 and 24. These drawings are the same except for the shading used to illustrate the contribution of the five levels of process masks used in diode fabrication. The anode of the diode is the small metal finger with an external connection in the center of the drawing for level D. The cathode is the larger surrounding metal.

FIGURE 23. Type A Diode (3 Square Micron)—N- and N+ Implants.




The contact area of the diode is that portion of the anode finger which is over the implant field. This is a rectangular area 0.8 microns by 3.6 microns in size. Total area is approximately 3 square microns. There is an ohmic contact on each side of the anode finger. This design is aimed at a configuration with minimum diode resistance per unit area.

FIGURE 24. Type A Diode—Metal Layers and Recess Insulator.



A type B diode with a 4.0 square micron diode contact area is illustrated in Figure 25. These dimensions are considerably more relaxed than the dimensions for the type A diode. The anode finger for this configuration is the right metal path in the drawing for level D. The cathode is the left metal path. The area of the diode in this configuration is 4.0 microns by 1.0 micron.

FIGURE 25. Design of a 4 Square Micron Type B Diode.




The implant areas are identical in this design as they were in the design of the type A diode. The advantage of the type B diode is the better process yield and the greater diode area in a given amount of package real estate.

The type A and type B diodes have differing electrical characteristics as a result of the difference in resistance to the ohmic contact. Plots of the forward voltage drop versus forward current are shown in the graph in Figure 26. The type B diode curve is approaching a straight line toward the right of the graph. This is because the 650 ohm forward resistance of the diode is dominating the electrical characteristics in this current range.

The type A diode has a much steeper curve in the voltage and current ranges shown. The type A diode curve is just beginning to straighten at the top of the graph. This slope represents the forward resistance of 200 ohms. Both curves reflect an underlying exponential component which becomes resistive when the current is high enough.

FIGURE 26. Diode Electrical Characteristics.




Another attribute of the diodes of interest in circuit design is the reverse biased capacitance. This is the loading factor, particularly for the cell input diodes, in switching state from one voltage level to another. This subject is addressed in the next section under the heading of capacitors. Similar information is applicable to the type A and B diodes.

2.5 Capacitor Design

Capacitors are formed from large contact area diodes. When used as a capacitor the potential across the diode is always reversed from that which would cause diode forward current. A design for a 95 square micron capacitor is illustrated in the drawings in Figures 27 through 30.

FIGURE 27. Design of a 95 Square Type C Diode/Capacitor.




Level A (N- Implant) and Level B (N+ Implant).

Most capacitors used in cell design are much larger than the 95 square micron capacitor used for the illustrations. Capacitors in actual cells are usually

several hundred square microns to several thousand square microns. However, the basic structure is the same as for the small capacitor illustrated.


FIGURE 28. Type C Diode/Capacitor—Ohmic Metal.



Level C: Ohmic Metal.

The capacitor illustrated can also be viewed as a diode. The drawing for level D (Figure 29) shows the external connections for the device. The cathode of the device, viewed as a diode, is the narrow metal, shaped like an upside down “T”. The anode is the large metal, shaped like an upside down “U”. The cathode is always positive with respect to the anode when used as a capacitor.

FIGURE 29. Type C Diode/Capacitor—First Layer Metal.




Level D: First Layer Metal.

The area for the capacitor is the area under the diode anode. This type C diode/capacitor has both N- and N+ implants. This causes the capacitance per unit area to be higher than if it were a single implant.

Notice that the recess insulator is only over the anode of all diodes.


FIGURE 30. Type C Diode/Capacitor—Recess Insulator.



Level H: Recess Insulator.


The series resistance of the implant area must be considered in capacitor design. The 200 ohms per square of resistance can become a design factor as the capacitor becomes large. This effect is minimized by breaking the capacitor contact area into smaller pieces and providing metal interconnections between the pieces. This design procedure is illustrated in Figure 31.

FIGURE 31. Capacitor Broken Into Two Pieces.



The electrical characteristics of the capacitor are graphed in Figure 32. The capacity per square micron on the vertical axis is the small signal incremental capacity for the given operating voltage. The voltages shown on the bottom of the graph are as viewed by a diode. That is, a positive voltage causes conduction and is generally not desired for a capacitor application.

FIGURE 32. Capacitor Electrical Characteristics.



The electrical characteristics of the capacitor are fairly linear in a reversed bias mode, but non-linear when forward biased. This characteristic allows for a capacitor which is more stable in the reversed bias mode. The capacitor's non-linear characteristics greatly affect the transistor and other diode voltage transition rates. It causes higher capacitance on the input diodes which in turn cause higher cell input loadings. It also retards the excursion rate on transistors when operated in the positive gate voltage region.

2.6 Resistor Design

A resistor is formed by placing ohmic contacts at the ends of an implanted area. In the CRAY-3 circuits resistors are made of N+ implant only. Some resistors are formed inside of a diode by increasing the N+ area from the cathode to the anode. The resistor value is determined by the number of squares of N+ area multiplied by the sheet resistance of the N+ implant plus the resistance in the connections to that area.

The width of the resistor is drawn 0.4 microns wider than needed to compensate for process shrinkage. Figures 33 and 34 show the layout of a 700 ohm N+ resistor. All resistors are drawn with a width of 4.4 microns or greater to ensure process controllability. Care must be used in designing an implanted resistor so that the resistor is not saturated. A saturated resistor will not have a linear electrical characteristic. Therefore, the resistor should be designed so that the maximum field strength applied will be 0.2 volt per micron.

FIGURE 33. Design of a 700 Ohm Resistor—N- and N+ Implants.






FIGURE 34. Ohmic and First Layer Metal in a 700 Ohm Resistor.



Level C: Ohmic Metal.



Level D: First Layer Metal.

2.7 Logic Package Design


A logic package in the CRAY-3 computer system is a GaAs integrated circuit die. The die is 3.835 mm square by 0.203 mm thick as it is placed on the printed circuit boards. It contains standard cell, electrical circuits to perform the computer logic functions and to communicate with other logic packages in the system. The average power dissipation of a logic package die is two to two and one-half watts. The level of integration represented by the package is equivalent to a 500 gate silicon emitter coupled logic (ECL) package. The speed of the circuits is equivalent to a 120 picosecond gate delay.

The CRAY-3 logic packages do not use emitter coupled logic (ECL). Neither do they use transistor gates of the type used in high-speed, silicon circuits. The circuits used in the CRAY-3 logic packages involve a level of diode logic followed by either two or three levels of inverting amplifiers. This is referred to as Schottky-Diode FET Logic or SDFL. A block diagram representation of two circuit types is shown in Figure 35.

A low voltage level of 500 millivolts with respect to ground is considered the static "one" state. A high voltage level of 1500 millivolts with respect to ground is considered the static "zero" state. The circuit threshold voltage is approximately 900 millivolts.

A logic package design is divided into a number of standard parts which are interconnected with metal paths to form the desired integrated circuit. The standard parts are called "cells" (or "macros") and are taken from a digital library. The current CRAY-3 Cell Library contains 36 standard cells and 33 routing and placement macros (these are described at the end of this chapter). The two cells illustrated in the block diagram are the D type cell and the E type cell. These are the two cell types which are used for logic generation in a package design where the outputs of the cells are used internally and do not leave the package. Cell types which drive transmission paths outside of the integrated circuit package are called F cells and will be described later.

FIGURE 35. Logic Package Block Diagrams.




The logical functions illustrated in the block diagram are called 5 by 3 logic. This means that each of the “and” circuits may have five input signals and there may be three such circuits with inputs to the common “or” connection. This is the average level of complexity in CRAY-3 logic cells. The maximum depth of logic allowed by the cell designs is 5 by 5. Each cell type has a number of “or” configurations which are defined in Table 7 at the end of this chapter.

The D type cells amplify and shape the signal resulting from the logic network. The output is inverted twice and is able to drive six active internal loads in the logic package. This cell type is the smallest of the types used in an integrated circuit package and occupies one cell site of the 128 available in a package.

The E type cells amplify and shape the signal resulting from the logic network and provide both polarities of output from the cell. This takes more power and twice the space of a D type cell but is necessary in generating complementary signals. Each output is able to drive six active internal loads in the logic package. The E type cells require two cell sites within the IC package.

The various logic cells used in an integrated circuit design receive power from a power grid structure which is the same in all package layouts with the exception of the VR package. This structure is illustrated in Figure 36.

FIGURE 36. Logic Package Power Grid and Cell Positions.




There are three power supply paths required for each cell in the IC package. The positive voltage (V_p) has a value of +3.3 volts with respect to ground. The negative voltage (V_n) has a value of -1.2 volts with respect to ground. The third path is the ground connection (V_g) for the reference potential.

The power distribution paths shown in Figure 36 are superimposed on each other and appear to be shorted. First layer metal carries V_g . Second layer metal distributes the V_p and V_n power to the cell positions. The V_p second layer metal completely covers the V_g power bus. This structure is over an implanted area which forms a large reversed bias diode/capacitor between the streets where the individual logic cells are located. These large capacitors are used to

supply current during transient operation. The Vn second layer metal goes down the middle of each street and is not shown in the illustration.

Sixteen of the 128 cell sites are illustrated in Figure 36 by the small rectangles on H street. Interconnections between package cells pass down the streets to the central avenue. There are 20 path positions for horizontal trace lines in each street and 56 path positions for vertical trace lines in the central avenue.

FIGURE 37. Logic Package Pinout.



Around the periphery of the logic package are 52 bonding pads. Ten of these pads connect to the power grid. Two pads near each corner are connected to Vp and Vg. Two pads on opposite sides are connected to Vn. Signal inputs to the logic package travel around the edge of the power grid and enter the central avenue both from the top and from the bottom. Pads 09 and 11 are dedicated for clock input signals. These are the only input signals to cross the power bus. Signal outputs from the logic package generally pass over the power grid from

the far ends of the streets and connect to bonding pads along the sides of the package.

2.7.1 D Cell Design

The D cell is the smallest logic cell used in a CRAY-3 integrated circuit package. It also uses the least amount of power. This cell type is used to form two levels of logic and to provide the re-shaped and amplified output to other package cells in the integrated circuit. Only the normal output of the logic network is provided.


A circuit schematic drawing for the DD cell is shown in Figure 38. The logic configuration illustrated has a 5-5-5-5 logic generation. There are different combinations of "ors" for different logical functions. Again, a complete listing of these combinations is provided in Table 7 at the end of this chapter.

The device sizes in the circuit schematic are shown in plain text numbers. For transistors this is the width in microns. For diodes it is the diode area in square microns. The letter following the device size for diodes indicates the diode configuration type. The numbers in italics are arbitrary node numbers for purposes of reference.

The inputs to the diode network are indicated by nodes 03 and 04 for the first "and" circuit. The maximum number of inputs available for a D cell is 20. This limit is imposed by the physical space at the street curb within the IC package.

The diode network at nodes 03, 04 and 05 makes the "and" function of each cell. Each input is connected to a 3A diode. The 12A diode is a short hand notation and represents 4 more 3A diodes with the cathodes all tied to node 05. Each "and" section contains 5 inputs. The input circuitry is designed to handle a voltage swing of 1.0 volt centered at 1.0 volt. The input diode along with the 700 ohm resistor and 4 micron current source tied to Vn causes about a 1.0 volt drop from the input to node 06. The input changes from a high value of 1.5 volts to a low value of 0.5 volt making the high value 0.5 volt and the low value -0.5 volt at node 06. The 8 micron transistor at node 06 begins to turn on when the voltage at node 06 goes above -0.5 volt. When the voltage at node 06 goes above 0 volts the transistor will over power the 8 micron FET at node 07 that was holding the voltage up and cause the voltage at node 07 to be switched from a high value of about 1.5 volts to a low value of 0.15 volt. The high value at node 07 is determined by the two diodes at nodes 07 and 08. These diodes clamp the voltage at node 07 to about 1.5 volts. The next transistor at nodes 01, 07 and 10 is a source follower. Whatever voltage is applied to the gate at node 07 will be produced at node 10 as long as no other circuitry is holding the voltage at node 10.

FIGURE 38. DD Cell Schematic Diagram.



The “or” section of the cell is at node 10. The 24 micron transistor at node 11 represents 3 more input structures all “ored” together at node 10. The DD cell is a 5 input “and” gate, “ored” 4 times at node 10 making a 5-5-5-5 logic function. Other logic configurations can be made by reducing the number of input sections tied to node 10. If any input on any one of the input sections has 1.5 volts (a logic “0”) applied to it, node 06 will be at a high value of 0.5 volt and node 07 will be at a low value of 0.15 volt while node 10 will be at a high value of about 0.4 volt. Node 10 is being held up by the other input sections tied together at node 10. If the logic configuration had less input sections tied together at node 10, the voltage at node 10 would be reduced to as low as 0.2 volt.

The 4B diode and 600 ohm resistor, with the 8 micron transistor used as a current source, causes the voltage at node 13 to be lowered by about 1.0 volt from node 10. This causes a high value of about 0.5 volt and a low value of about -0.55 volt at node 13. The 600 ohm resistor is reduced to 480 ohms in DA, DB and DC cells. These D cells have less “ors” tied at node 10. This allows the voltage at node 10 to be reduced when at its low voltage level. The 480 ohm resistor produces less voltage drop. This causes the voltage at node 13 to be about the same for all D cells. This voltage swing at node 13 allows the 8 micron transistor at nodes 14, 13 and 00 to operate like the 8 micron transistor at nodes 07, 06 and 00. When the voltage at node 13 is less than -0.5 volt the transistor is off, allowing the voltage at node 14 to be a high value of about 1.5 volts. When the voltage at node 13 starts rising, the 8 micron transistor starts to turn on. When the voltage at node 13 goes above 0 volts, the 8 micron transistor at nodes 01 and 14 is over powered. Then the voltage at node 14 is inverted from 1.5 volts to about 0.15 volt. Again, the two diodes at nodes 14 and 15 clamp the high value to about 1.5 volts at node 14.

The output transistors for the cell are the 64 micron wide FETs at node 17. The output transistor tied to V_p is a source follower. When the voltage at node 14 is at 1.5 volts, the voltage at node 17 will follow to 1.5 volts or greater. When the voltage at node 14 is at 0.15 volt, the voltage at node 17 goes to 0.5 volt or less. The output FET connected to ground acts as a current source to hold the voltage down at node 17. Without this 64 micron pull down FET the output levels would be much higher. The D cells were designed to drive up to 6 active inputs internal to the GaAs integrated circuit package. The propagation delay through the is approximately 300 picoseconds.


2.7.2 E Cell Design

The E cell is the most versatile of the cells. This cell provides both a normal and a complementary output signal. The E cell requires twice the power and twice the space of a D cell. One advantage of the larger space is that there is more street curb for input diodes. The maximum number of input diodes for the E cell is 25 as compared to the D cell maximum of 20. Each output signal from the E cell can drive six active loads internal to the integrated circuit package.

A circuit schematic for the ED cell is shown in Figure 39. The E cell is very similar to the D cell. The input section of the E cell is identical to the D cell except for the transistor at nodes 01, 07 and 10. This transistor is twice the size as the one in the D cell in order to drive two sets of level shifting circuitry connected to node 10. One path drives one inverting amplifier which is identical to the D cell. The other path drives two inverting amplifiers in series to produce the complement output. The two paths act the same except for the circuitry from node 19 to node 25. This is the extra inverter path.

The voltage at node 19 has a high value of about 0.5 volt and a low value of about -0.55 volt. This is the same as at node 13. The voltage at node 10 increases as more "ors" are connected together at node 10 just like the D cells. The 600 ohm resistor at nodes 13 and 19 is reduced to 480 ohms for E cells with less "ors" tied together at node 10. This provides the same voltage swing at nodes 13 and 19 for all E cells. When the voltage at node 19 is below -0.5 volt the 8 micron wide FET connected to ground is turned off. This allows the 8 micron wide FET connected to V_p and node 20 to pull node 20 to about 1.5 volts. The two diodes at nodes 20 and 21 clamp the high value to 1.5 volts. When the voltage applied to node 19 rises, the 8 micron transistor begins to turn on. When the voltage goes above 0 volts, node 20 will begin to be pulled down to 0.15 volt.

FIGURE 39. ED Cell Circuit Schematic.



The 8 micron source follower at node 20 causes the voltage levels at node 23 to be about 0.2 volt for its low value and about 1.5 volts for its high value. The 4B diode, 600 ohm resistor and the 8 micron current source drop the voltage about 1.0 volt from node 23 to node 25. This voltage excursion is the same for all E cells no matter what the logic configuration. The voltage at node 25 changes from a high value of about 0.5 volt to a low value of about -0.75 volt. The output stage of the complement output works the same way as the output stage of the normal output (identical to the D cell output stage).


The circuit delay through the two paths of the E cell is not the same. The transit time through the normal output path is about 350 picoseconds whereas the transit time through the complementary output path is about 400 picoseconds for the output falling edge and about 500 picoseconds for the output rising edge. Therefore, care must be taken in logic package design to avoid successive E cell paths with too much loading.

2.7.3 F Cell Design

The F cell is used to drive signal lines leaving the integrated circuit package and traveling to other GaAs packages in the system. The structure of the F cell is similar to the structure of the D cell. The difference is in the size of the devices at the “or” connection and the output stage. A circuit schematic for the FE cell is shown in Figure 40. The F cell requires twice the space and more than twice the power of a D cell because of the larger components involved.

The F cell drives a transmission line. This transmission line is terminated with a 55 ohm resistor located in the resistor plate (a few signals are terminated with 55 ohm resistors located in a special IC package). Without the terminating resistor the F cell output levels would be too high and would never go low. The output of the F cell can also be used to drive signals internal to the die. The output voltage levels of the F cell are slightly compressed compared to the D or E cell output voltage levels. This is the effect of the terminating resistor instead of the current source pull down transistor in the D and E cells.


FIGURE 40. FE Cell Circuit Schematic.



Node 00 - Circuit Ground
Node 01 - Positive Power Supply (VP)

Improved versions of the FA through FE cells were added to the CRAY-3 Cell Library in 1992. These were designated as FM through FQ. These new F cells were designed to improve the output levels and provide a buffered output signal for use within the integrated circuit package. The schematic diagram for the FQ cell is shown in Figure 41.

FIGURE 41. FQ Cell Circuit Schematic.




The new F cells are FAST tested (see the section on FAST testing in Chapter 5) for a Voh (voltage output high) of 1.35 volts and a Vol (voltage output low) of 0.45 volts, whereas the FA through FE series were FAST tested for a Voh of 1.30 volts and a Vol of 0.58 volts. As you can see, the greater improvement was accomplished at the lower end of the voltage swing.

If you take a few moments to compare the schematic diagrams of the FE cell and the FQ cell, you will notice that the cells are basically the same from the inputs all the way to nodes 13 (in the FE cell) and 130 (in the FQ cell). The only major differences between the cells are shown by the shaded areas in Figure 41. These include the level shift circuitry between nodes 130 and 190, the dynamic level shift circuitry to improve the output swings of the output stage, and the buffered output circuitry.

2.7.4 FF and FG Cell Design

The FF and FG cells are used to drive signal lines leaving the GaAs integrated circuit packages and traveling to a silicon package in the system. A circuit schematic of the FF cell is shown in Figure 42. The cells are identical to the standard F cells except for the removal of the clamping diodes at node 14 in the FF and FG cells. The absence of these clamping diodes allows the output high value (V_{oh}) to be about 2.25 volts instead of 1.5 volts. This higher voltage swing is needed to interface properly with the silicon memories being used in the CRAY-3. The termination of the FF and FG cell is a 55 ohm resistor. This is used to reduce the power and allow for an easier transition to the higher voltage level. Because of the higher output level of the cell it can not be used to drive any GaAs inputs used in the CRAY-3 unless its output level is shifted lower through an AG cell.


FIGURE 42. FF Cell Circuit Schematic.



2.7.5 FJ and FK Cell Design

The FJ and FK cells are improved versions of the FF and FG cell designs. They were designed first, with the FM through FQ series of cells following sometime later. Like the FM through FQ cells, the FJ and FK cells have improved output levels to more reliably interface with silicon parts. They are FAST tested at Voh levels of 2.5 volts and Vol levels of 0.35 volts (see the section on FAST Testing in Chapter 5). A schematic diagram of the FJ cell is shown in Figure 43.

FIGURE 43. FJ Cell Circuit Schematic.



2.8 Clock Amplifier Design

Each logic package has a clock signal amplifier in addition to the logic package cells. The clock amplifier receives differential external signals from bonding pads 09 and 11. These bonding pads are assigned this function in all logic package layouts. Pad 11 provides the normal clock signal input and pad 09 provides the inverted clock signal input.

The clock amplifier is located in D street west for all logic package layouts, with the exception of the Vector Register package. The differential input signals from bonding pads 09 and 11 pass over the power bus and enter the clock amplifier circuits from the far west end of D street. The clock amplifier consists of two identical cell designs. These cell designs are designated as CA in the logic package cell library. Each CA cell requires two D cell positions.

The circuit schematic in Figure 44 shows both CA cells. One cell is located on each side of D street west. The points in the schematic labeled 05 and 06 cross the street and connect to the equivalent points in the circuit on the opposite side of the street. Again, the plain text numbers are device sizes and the italic numbers are node designations for purposes of identification. The clock amplifier circuits invert two times in each of the cell configurations. The first stage of the amplifier is differentially coupled with the opposite cell. The last stage is an inverting amplifier.

The inputs to the clock amplifier are AC coupled by a 3000 type C diode/capacitor reverse biased at nodes 03 and 13 to nodes 04 and 14 respectively. Node 05 is biased to a constant voltage of about 0.72 volt by an 8 micron transistor used as a current source and an 8B diode in the input section of both CA cells that are used for the normal and complement clocks. Inputs to the clock amplifier are 180° out-of-phase. This phase shift is kept through the two CA cells used to produce and distribute the clock signals to each logic cell.

The 4K ohm resistors from node 05 to nodes 04 and 13 supply the reference bias voltage for the AC coupled inputs. This causes the voltage at node 04 and 14 to be centered at about 0.72 volt with a swing of about 0.5 volt when the inputs are driven by clock relay circuits (CB cells). The RC time constant produced by the 3000 type C diode/capacitor and 4K ohm resistor is several times larger than the 500 mHz input signal reaction time. A shift other than 180°, a frequency change for a short period of time, or a change in amplitude to the inputs of the clock will be stabilized within the clock amplifier because of this large RC time constant.

[illegible]

Node 05 and node 06 are tied together between the two CA clock amplifiers to give the stability already described. The circuitry after node 07 and node 15 of the respective CA cells are independent of each other and function identically, but 180° out-of-phase to each other. Only the first CA cell will be discussed in the rest of this clock amplifier description.

3207 - CRAY-3 Hardware Description Manual


volts the voltage at node 08 is at -0.6 volt. The large RC time constant here has a similar effect as the one in the input section.

The last stage of the clock amplifier causes another inversion and operates similarly to the D cell output stage. The transistors on the output stage of the clock amplifier are sized to allow the falling propagation delay to be faster than the rising propagation delay. This asymmetry causes the crossover point of the normal and complement clock to be about 0.5 volts or lower instead of 1.0 volt. This allows the latch used in the GaAs logic packages in the CRAY-3 to operate with more margin on the closure path of the latch. The clock relay circuit (CB cell) works similarly to the clock amplifier circuit except that its output stage produces a crossover point at 1.0 volt (symmetrical).

The clock amplifier outputs are designed to drive up to 60 internal inputs. The normal and complement clock amplifier outputs are balanced within each GaAs logic package by routing their outputs to ESD diodes for capacitive loading. ESD circuits are added to the clock outputs so the total, after route load is within 75 femtofarads of 3000 femtofarads for each of the clock output signals. The normal and invert clock output signal are also balanced within 75 femtofarads of each other. This causes minimum clock skew within the die and from die to die.

Another clock amplifier has been added to the CRAY-3 cell library. It is the CC cell. This cell replaces the CA clock amplifier in all IC packages currently used (except for the VR and VV packages). The CC cell is similar to the CA and CB cells except for additional feedback circuitry on the last stage. This causes greater asymmetry in the clock crossover point. A circuit schematic showing both CC cells is found in Figure 45.

FIGURE 45. CC Cell Circuit Schematic.




A 40 micron transistor, 4k ohm resistor and 3000 type C diode/capacitor were added to each clock amplifier to make the CC cell. This requires each CC cell to occupy three D cell sites. This inverting feedback path causes the falling edge of one clock signal to make the rising edge of the other clock signal occur about 250 pico seconds later. This produces clock pulses which are asymmetric with each other. The pulse low voltage time is longer than the pulse high voltage time. This helps the latch hold the output data before the input data is released on negative pulses.

The increased asymmetry of the CC clock amplifier entails one drawback in its operation. It does reduce latch closure problems, but low-to-high pulses (negative) from the latch will be longer while high-to-low pulses (positive) will be shorter. Generally, this is not a problem since latch-to-latch operation has about one nanosecond of cycle time. However, in some logic packages where there are long delay chains, the cells that pass on the signal in the chain can be too close to each other (the signal is passed on too quickly) causing a latch to fail. This situation is now taken into account by the cell layout programs in Artwork and corrected when the cells are placed.

2.9 Other Circuits

The CRAY-3 cell library consists of more than clock and logic cells. Electrostatic discharge (ESD) protection and power supply surge protection is performed in the diode structures of the AF cell. A circuit schematic for the AF cell is shown in Figure 46.


FIGURE 46. AF Cell Circuit Schematic.



The ESD diodes are connected to all inputs of the GaAs die. This allows voltage excursions 0.7 volt higher than V_p and 0.7 volt lower than V_g to be shunted away from the inputs. This voltage excursion, if received by the input structures of the logic cells, could be destructive. Power surges above 3.5 volts and below -1.4 volts will be absorbed by the diode circuitry in the AF cell. The AF cell requires one cell site. It is placed in all cell sites not used by other cells.

The AG cell also requires one cell site. Its circuit schematic is shown in Figure 47.


FIGURE 47. AG Cell Circuit Schematic.



The AG cell is used to translate the voltage levels from the 4K by 1 CMOS SRAM to the GaAs logic levels. The resistor and diode structure of the AG cell performs this function and also provides termination and ESD protection. The 400 ohm resistor produces a voltage drop when the input signal is above 1.5 volts. The two 32B diodes-to-ground clamp the high value to about 1.5 volts when the input voltage is at 3.0 volts. When the input signal is about 1.5 volts to -0.7 volt the current through the 400 ohm resistor is negligible. This allows the output to trace the input in that voltage range. The other 32B diode which is reversed biased to ground provides protection from voltages below -0.7 volt.

The AH cell is used to translate the signals from other parts with higher voltage excursions to the voltage levels required by the logic circuits of the CRAY-3. The 256K by 1 CMOS SRAM high value voltage level of 5.0 volts is clamped to about 1.5 volts by the AH cell. The AH cell requires one cell site. Its circuit schematic is shown in Figure 48.

FIGURE 48. AH Cell and AI Cell Circuit Schematics.




The two 800 ohm resistors divide the voltage in half when the input voltage is 3.0 to 1.5 volts. This produces a 1.0 volt output to the logic cells when 2.0 volts is at the input with only 1.25 mA. The 32B diodes-to-ground provide voltage clamping for protection.

The AI cell is a 55 ohm resistor used to terminate the signals from the FF/FJ and FG/FK cells to the silicon memories. The cell contains four N+ only implanted 220 ohm resistors in parallel. It uses one cell site.

2.10 The Vector Register Package

The Vector Register package used in the CRAY-3 merits special attention because of its unique design, complexity and extensive usage. The VR package has the highest component density of all the GaAs packages used. It is used over 3,000 times in a 16-processor CRAY-3.

FIGURE 49. Vector Register Block Diagram.



The VR package contains a 256-bit memory with a 64 element address and four data bits. Six address bits are delivered to the package for readout

selection and six are delivered for write selection. Write one and write zero signals are supplied for each data bit as can be seen in Figure 49. Input signals are latched on an even clock phase (see Figure 50) and readout signals are latched on the following odd clock phase (see Figure 51).

FIGURE 50. VR Write Timing.





FIGURE 51. VR Read Timing.




2.11 The Silicon Memory Packages.

The CRAY-3 uses four types of silicon integrated circuit packages for its memory requirements. These include the SC packages used for local memory and the SD, SF or SM packages used for Common Memory.

The SC die are 4,096 by 1 static memory circuits used for local memory. The die are used in pairs (mounted side by side) and are treated as a single unit. The circuit has a six nanosecond cycle time for both read and write functions.

FIGURE 52. Size Comparisons of the Memory Die.




The SD die are 256K by 1 static memory circuits used for Common Memory. Four packages are used in parallel to provide 1,024K bits of memory per bit position in each Common Memory bank. The circuit has a 34 nanosecond write cycle and a 34 nanosecond read cycle. The data is sampled at 24 nanoseconds into the read cycle.

The SF die are one meg by one static CMOS memory circuits used for Common Memory. Two packages are used in parallel to provide 2,048K bits of memory per bit position in each Common Memory bank. The circuit has a 25 nanosecond cycle time for both read and write functions. Again, the data is sampled 24 nanoseconds into the read cycle.

The SM die are one meg by four static CMOS memory circuits which are being planned as a memory enhancement to the CRAY-3 as this revision of the Hardware Description Manual goes to press. The use of these die would boost the Common Memory capacity of the CRAY-3 four times per octant of memory over the standard SD die. The SM die have an address cycle time of 25 nanoseconds for both read and write. If the SM die implementation is successful, the SF die will not be used.

Pinout diagrams for the memory die can be seen in Figures 53 through 55.

FIGURE 53. SC Die Pinout.



Two SC Die shown as a single unit.

FIGURE 54. SD and SE Die Pinouts.





FIGURE 55. SM Die Pinouts.



2.12 The CRAY-3 Cell Library.

The integrated circuit packages (with the exception of the VR package discussed above) used in the CRAY-3 are created by using different combinations of cells from the cell library. This library is a repository of the digitized data which was created when these circuits were designed. The digitized data was created in one of two ways:

- Directly. The digitized data can be created directly by calculating the mathematical coordinates for the layout of all the rectangles used to produce each circuit. The resulting numbers are then typed in the proper format which can be read by the computer programs used in the Artwork Department to create the IC packages.
- Indirectly. The digitized data can be an output file of a Computer Aided Design program. For the CRAY-3 this would include either CrayDraw (an in-house CAD program) or Cadence (a commercial design tool).

The following table lists all of the circuits currently part of the CRAY-3 Cell Library:

TABLE 7. The CRAY-3 Cell Library.

Cell Name	Size (in grid units)	Type	Inputs	Outputs
AA	38100 x 38100	Final Assembly. Calls and places all other cells. Inscribes the IC package name.	Does not apply.	Does not apply.
AB	19052 x 19052	Quadrant Assembly. Calls and places one quadrant of the crack stoppers and the power grid.	"	"
AC	1380 x 1380	Bonding Pads.	"	"
AD	152 x 3152	Vertical Crack Stoppers.	"	"
AE	3152 x 152	Horizontal Crack Stoppers.	Does not apply.	Does not apply.
AF	2400 x 648	Null Cell. Placed in empty cell locations. Provides power surge and ESD protection.	Two ESD. One power surge.	None.
AG	2400 x 648	SC Readout Interface. Includes one 400 ohm resistor. Reduces the output voltage from SC die and TTL circuitry to match the lower input level of GaAs die.	One.	One.

Cell Name	Size (In grid units)	Type	Inputs	Outputs
AH	2400 x 648	SD Readout Interface. Includes two 400 ohm resistors. Reduces the output voltage from SD die to match the input level of GaAs die.	One.	One.
AI	2400 x 648	Termination Resistor. Includes four 220 ohm resistors in parallel (55 ohms).	One.	None.
AJ	19052 x 19052	Quadrant Assembly for use with the FJ and FK cells. Allows the wider H macros to cross over the power grid.	Does not apply.	Does not apply.
CA	4800 x 648	Clock Amplifier. Clock signal distribution internally. 150 ps clock overlap. Used in the VR and VV packages.	One.	One.
CB	4800 x 648	External Clock Amplifier. Distributes the clock signal to other packages. Symmetrical.	One.	One.
CC	7200 x 648	Clock Amplifier. Clock signal distribution internally. 250 ps overlap. Used in all clocked ICs except for VR, VV, and ZT.	One.	One.
DA	2400 x 648	Logic Cell. The D and E series logic cells can feed up to 6 active internal loads at a time.	5	True.
DB	2400 x 648	Logic Cell.	5-5	True.
DC	2400 x 648	Logic Cell.	5-5-5	True.
DD	2400 x 648	Logic Cell.	5-5-5-5	True.
EA	4800 x 648	Logic Cell.	5	True and False.
EB	4800 x 648	Logic Cell.	5-5	True and False.
EC	4800 x 648	Logic Cell.	5-5-5	True and False.
ED	4800 x 648	Logic Cell.	5-5-5-5	True and False.
EE	4800 x 648	Logic Cell.	5-5-5-5-5	True and False.

Cell Name	Size (in grid units)	Type	Inputs	Outputs
EH	4800 x 648	Logic Cell. EH through EL are modified versions of EA through EE. Whereas the EA series has three inversions in the complement path, the EH series has only one inversion to increase the speed of the complement path.	5	True and False.
EI	4800 x 648	Logic Cell.	5-5	True and False.
EJ	4800 x 648	Logic Cell.	5-5-5	True and False.
EK	4800 x 648	Logic Cell.	5-5-5-5	True and False.
EL	4800 x 648	Logic Cell.	5-5-5-5-5	True and False.
FA	4800 x 648	Logic Cell Transmitter. The F series logic cells feed GaAs loads external to the package.	5	One.
FB	4800 x 648	Logic Cell Transmitter.	5-5	One.
FC	4800 x 648	Logic Cell Transmitter.	5-5-5	One.
FD	4800 x 648	Logic Cell Transmitter.	5-5-5-5	One.
FE	4800 x 648	Logic Cell Transmitter.	5-5-5-5-5	One.
FF	4800 x 648	Signal Driver Interface. Provides wider voltage swings for feeding silicon circuits. Vol 0.5 volts. Voh 2.4 volts.	5	One.
FG	4800 x 648	Signal Driver Interface.	5-5	One.
FJ	4800 x 648	Improved version of FF. Vol 0.25 volts. Voh 2.4 volts.	5	One.
FK	4800 x 648	Improved version of FG.	5-5	One.
FM	4800 x 648	Improved version of FA.	5	One external and one internal.
FN	4800 x 648	Improved version of FB.	5-5	One ext. One int.

Cell Name	Size (In grid units)	Type	Inputs	Outputs
FO	4800 x 648.	Improved version of FC.	5-5-5	One ext. One int.
FP	4800 x 648.	Improved version of FD.	5-5-5-5	One ext. One int.
FQ	4800 x 648.	Improved version of FE.	5-5-5-5-5	One ext. One int.
GA	7640 x 7428	G series macros are used to connect the outputs of F cells to bonding pads. They are not used with the FJ and FK cells. H street to Pad 52.	Does not apply.	Does not apply.
GB	7640 x 7676	H street to Pad 02.	"	"
GC	7640 x 10728	G street to Pad 01.	"	"
GD	7640 x 10976	G Street to Pad 03.	"	"
GE	7640 x 17328	E Street to Pad 05.	"	"
GF	7640 x 17576	E Street to Pad 01.	"	"
GG	7640 x 7428	H Street to Pad 01.	"	"
GH	7640 x 7676	H Street to Pad 03.	"	"
GI	7640 x 10748	G Street to Pad 05.	"	"
GJ	7640 x 15828	G Street to Pad 07.	"	"
GK	17488 x 5856	H Street to Pad 51.	"	"
GL	17712 x 5856	H Street to Pad 50.	"	"
GM	17936 x 5856	G Street to Pad 49.	"	"
GN	18160 x 5856	G Street to Pad 48.	"	"
HA		The H series macros are about 10 microns wider than the G series to accommodate the greater current needs of the FJ and FK cell outputs. H Street to Pad 52.	"	"
HB		H Street to Pad 02.	"	"
HC		G Street to Pad 01.	"	"

Cell Name	Size (in grid units)	Type	Inputs	Outputs
HD		G Street to Pad 03.	"	"
HE		E Street to Pad 05.	"	"
HF		E Street to Pad 07.	"	"
HG		H Street to Pad 01.	"	"
HH		H Street to Pad 03.	"	"
HI		G Street to Pad 05.	"	"
HJ		G Street to Pad 07.	"	"
HK		H Street to Pad 51.	"	"
HL		H Street to Pad 50.	"	"
HM		G Street to Pad 49.	"	"
HN		G Street to Pad 48.	"	"

In the next chapter we will examine more closely how these cells are used to create all of the 498 different types of GaAs Integrated Circuit packages used in the CRAY-3.

Integrated Circuit Artwork

In this chapter we turn our attention to the basic building blocks of the CRAY-3: the integrated circuit packages. Here we learn how the software programs in the Artwork Department use the individual pieces (the cells) from the CRAY-3 Cell Library to create each integrated circuit package.


3.1 The IC Assembly

The software programs used to create the integrated circuits begin with a generic IC layout and automatically customize that layout to conform to the specifications found in the IC Boolean (or Cover Sheet) for each unique IC package. The IC Boolean is created by the logic design engineers and specifies the logical functions that each integrated circuit package must perform. This allows the logic design engineers to easily create every individual integrated circuit package for a specific application in the overall design of the machine.

The generic IC assembly is the same for all but four of the 498 IC types that are used in the CRAY-3. What makes each of the 494 ICs that use the generic layout unique is the different combinations of smaller pieces and routed lines that are added to the generic assembly. The generic assembly is analogous to a sub-division plan designed by city engineers. The streets and avenues have been laid out and the size of each available real estate lot along those thoroughfares has been set. However, the overall complexion of the sub-division will depend on which lots have homes built on them and what

type of homes are built. These choices will also affect the traffic flow patterns on the streets and avenues.

FIGURE 56. Generic IC Assembly.



There are eight streets and one avenue in each generic IC assembly. The avenue runs north and south down the center of the IC, so each of the eight streets are divided into a west side and an east side (see Figure 56). Each half of each street has four cell positions on the north side of the street and four cell

positions on the south side of the street where different types of logic cells can be placed depending on the overall needs of each particular IC package.

Each CRAY-3 IC package has 52 bonding pads around the periphery of the package. These pads are used for signal passage into and out of the package. Of the 52 pads there can be a maximum of 34 logic signal inputs with six outputs, or a maximum of 38 logic signal outputs with two inputs. Pads 9 and 11 are reserved for clock signal inputs, and pads 4, 6, 8, 10, 12, 30, 32, 34, 36 and 38 are reserved for power and ground connections (refer to Figure 37 in Chapter 2, page 51).


The manner in which the generic IC assembly is customized to make it into a specific IC package is specified by the IC Boolean (or Cover Sheet) received from a logic design engineer on a micro-diskette. This document, which is about four to eight pages in length, tells the PlaceCells and Route programs which logic cells need to be placed along the streets, where to place them, and the pads and cells that need to be inter-connected with trace lines to enable the proper combination of signals to get around within the IC. Through its Pad Assignment section the IC Boolean document also informs the computer-router which signals need to be received and sent out to trace lines on the circuit boards through a gold pin bonded to the bonding pads.

When the Artwork Department receives the IC Boolean for a particular IC package from a logic design engineer the data is first pre-tested for boolean and syntax errors. Once the IC package data has passed pre-testing it can then be routed, tested, fractured and sent to glass where glass masks will be made of each layer. These glass masks are like large glass slides of the drawn data which are used to project the image of the IC parts for each layer onto the GaAs wafers during the manufacturing process at Cray Computer Corporation's GaAs Fabrication Facility.

To understand how each generic IC assembly is customized into a specific IC package one needs to have a basic understanding of how the text data represents the digitized data which a computer and its associated software can understand and use. The key to this understanding is in understanding IC macros.

IC macros (sometimes called cells), as used in the development of the CRAY-3, are small text files which tell a CAD program where and how to draw rectangles on several different drawing layers using a pre-determined grid size. Each macro line has five parts as illustrated in Figure 57. Each grid unit is equal to one-tenth micron.

FIGURE 57. Parts of a Macro Text Line.



Macros can be nested. That is, one or more macros can nest, or be in side of other macros. In other words, one macro can call up (reference) another macro by giving the CAD program the name of the macro and the x and y coordinates of where to place it. Each of the 494 CRAY-3 ICs which use the generic IC assembly use a master macro, called AA, which calls up all the other macros needed to customize the AA macro into each particular type of IC package. The generic IC assembly is made up of six macros as shown in Figure 58. Macros AB, AC, and AF are all nested within the AA macro. In other words, the AA macro calls up those macros and places them at precise x, y coordinates within itself. Macros AD and AE are nested within the AB macro.

FIGURE 58. Macros and the Generic IC Assembly.

AA	Calls the other macros and places them at the correct grid location. Contains the macros for the bonding targets. Contains the macros for the saw streets. Contains macros for routing clock inputs from pads 9 and 11.
AB	The Quadrant Assembly. This is basically one-quarter of the power grid. The AA macro calls this macro 4 times to make up a full assembly.
AC	This macro draws one bonding pad.
AD	This macro draws the vertical crack stoppers for the AB macro.
AE	This macro draws the horizontal crack stoppers for the AB macro.
AF	This macro draws a null cell which is placed at every empty cell position along the streets.

Since the grid units that are used in the design of the ICs are equal to one-tenth of a micron, the smallest object that could be drawn in the ICs themselves would have to be at least one-tenth micron on a side. The overall size of the IC package, as it is drawn, is 38,100 by 38,100 grid units or 3,810 microns on a side (3.810 mm square).

At this point one may wonder how complex ICs are drawn with such simple lines of text. But it is much easier to understand when one realizes that the integrated circuits contain nothing but rectangles of varying sizes and on nine different layers. When one looks at a finished die under low magnification the bonding pads appear to be round dots of gold. However, as you increase the magnification you see the "round" pads begin to appear more and more like a group of various size rectangles joined together.

Now that we have a basic understanding of what a macro is like we can proceed to show how the generic IC assembly is created by a computer using macros designed in simple text format. We can then go on to show how the generic assembly is customized to make a real IC package.

In order to do this we will reproduce several portions of the generic AA macro and then illustrate what those macro lines actually do. During this process we will also be illustrating the basic principles used in a text macro to draw the CRAY-3 integrated circuits.

FIGURE 59. First Part of the AA Macro.

Macro AA - Final Assembly


Macro size 38100 x 38100 - Revision Zero - June 22, 1990

0 0 AB +19050 +19050.	lower left frame
0 38100 AB +19050 -19050.	upper left frame
38100 38100 AB -19050 -19050.	upper right frame
38100 0 AB -19050 +19050.	lower right frame

In the first illustration of the AA macro, shown in Figure 59, the macro is first named and identified. Then the overall size is given in grid units, the revision level is stated and the date of the last revision given. Notice that the overall size of the AA macro is 38,100 by 38,100 grid units, which is the overall size of a CRAY-3 GaAs die as drawn.

The next four lines of text specify the starting coordinates and the x, y directions for the placement of the four AB macros within the AA macro. This illustrates the nesting of one macro within another macro. Line one is telling the computer to call macro AB and place its starting coordinates at x, y coordinates of 0, 0 in the AA macro. At those coordinates in the AA macro the AB macro will be drawn in a positive x direction of 19,050 grid units (+19050) and in a positive y direction of 19,050 grid units (+19050). Notice that 19,050 is half of 38,100. This makes sense because the AB macro is one-quarter the overall size of the complete IC, so it would be half the length of one side of the IC. This portion of the AA macro, then, would place four AB macros in their proper place and orientation within the AA macro (which is the overall IC). Figure 60 gives a visual representation of the placement of the AB macro within the IC to help you better understand the process.

FIGURE 60. Placement of the AB Macro.



The next portion of the AA macro which we will examine is that part which calls up the AC macro (the bonding pad), placing it fifty-two times at the correct locations around the perimeter of the IC package. Only four of the fifty-two calls are shown in Figure 61.

FIGURE 61. AC Macro Calls.


580 580 AC +1380 +1380.	pad 01
+2540 +2540 AC +1380 +1380.	pad 02
-2540 +2540 AC +1380 +1380.	pad 03
+2540 +2540 AC +1380 +1380.	pad 04

The first line of that portion of the AA macro shown in Figure 58 calls the AC macro and places it at x, y coordinates of 580 and 580, respectively, to make up bonding pad one. At this location the macro is drawn positively in both the x and y directions. In the next line, for pad two, we are introduced to another principle in the use of macros to make up CRAY-3 ICs. Up to this point only absolute x, y coordinates have been shown. Now the + and - signs are used on numbers which are to the left of the letter designations (+2540 and +2540). These are not absolute coordinates, but are referenced to the preceding line. In other words, the +2540 for the x and y coordinates means that the AC macro for pad two will have a starting coordinate of 3120 for x and 3120 for y (580, the coordinate in the preceding line, plus 2540). It will then be drawn from that

new starting point in a positive direction for both the x axis and the y axis. In the third line, for bonding pad three, the coordinate values reference the preceding line again, since they are signed. They do not go back and reference the last mentioned absolute coordinate (580, 580), but reference the newly created coordinate in the preceding line. So the new starting coordinates for pad three would be 580 for x and 5660 for y (3120 minus 2540, and 3120 plus 2540).

Notice that with pad three the location has returned to the same horizontal position (x coordinate) as pad one, but gone higher in vertical positioning (y coordinate) as illustrated in Figure 62. By comparing the macro text shown in Figure 61 with the positioning of the AC macros illustrated in Figure 62 you should be able to figure out the relationships between the macro text and the placement of the macros in the IC package itself.

FIGURE 62. Placement of the AC Macros.



In the next example, shown in Figure 63, we encounter another principle used in IC macros. This is the drawing layer designation. Up to this point we have seen only two-letter designations representing the names of macros to be called, or referenced, by another macro. Notice that in this example we find only a single letter on each line. Single letter designations refer to drawing layers A through H, and R.


FIGURE 63. The Macro Lines for the Bonding Targets.

```

34208 600 D +160 +2000.          bonding targets
33288 1520 D +2000 +160.

3892 37500 D -160 -2000.
4812 36580 D -2000 -160.

```



The four macro lines shown in Figure 63 draw the two bonding targets for each IC. The bonding target drawn by the first two lines, which is the target which appears at the bottom, right corner of the IC package, has been illustrated in Figure 63. Each line tells the computer to draw a rectangle 160 by 2000 grid units on layer D at a specific beginning coordinate. The software knows that a rectangle is intended, so it automatically closes the rectangle even though only two sides are given in the data.


All of the basic principles involved in using integrated circuit macros have now been presented. One thing that may not be evident, however, is that macros are transparent to each other. In other words, one macro can specify a rectangle whose boundaries crossover into another rectangle specified by a completely different macro. A simple illustration of this is the two rectangles making up the bonding targets. They overlap each other at the center. But since both rectangles are on the same layer and will be filled in as a solid, plus mark in manufacturing, this creates no problems.

3.2 Routing an IC

Nearly all the IC packages that are used in the CRAY-3 use the generic IC assembly which is created using macros AA, AB, AC, AD, AE and AF. To create an actual working IC the computer uses the routing software to add the necessary additional pieces to the generic assembly (or final assembly) to make that specific IC. Once all the pieces have been placed at their proper locations within the AA macro, the computer connects all the pieces that need to communicate with each other using more macros, which it automatically creates, to make the routed trace lines in the streets and avenues. The computer knows how to customize each AA macro by reading the IC Boolean for a particular type of IC package.


FIGURE 64. Creating a CRAY-3 IC.

Step One.




The computer uses macros AA, AB, AC, AD, AE, and AF to create the generic IC assembly. It copies all of the needed macros from the current cell (macro) library stored on the hard drive.

Step Two.



The IC Boolean (or Cover Sheet) is read by the computer. From this data it knows what additional parts (macros) are needed to make up the specific IC package being specified by the IC Boolean. It copies these additional macros from the cell library placing them in the proper locations within the generic IC assembly (the AA macro).


Step Three.



The computer reads the pad assignments and boolean statements contained in the IC Boolean and creates the required macros needed to interconnect the parts with trace lines. It uses G or H macros from the cell library for connecting large output traces to the required pads.

In the next few pages we will expand on these three basic steps involved in creating an IC by using illustrations from the IC Boolean itself as well as illustrations from the IC packages being created. Figure 65 illustrates the first step: the creation of the generic IC assembly.


FIGURE 65. Step One: The Generic IC Assembly Is Created.



As soon as the generic assembly is drawn by the Route software, the Route program reads the IC Boolean and begins customizing the generic assembly into a specific IC package type. It does this by first reading the boolean statements to find out what types of logic cells are needed. It then reads the boolean placement to find out where to place the output cells. It automatically determines the proper placement of the remaining logic cells based on pre-determined guidelines (note the importance of cell placement in long delay

chains discussed in Chapter 2, page 64). Figure 66 shows the completed generic IC assembly with the logic cells and clock amplifiers for IC package EI placed along the streets in their proper positions. Notice that at this stage there are no routed trace lines in the avenue or the streets. The only signal trace lines present are the two input traces for the clock signals coming in from pads 9 and 11. These were drawn by the AA macro in step one.

FIGURE 66. Step Two: Placing the Required Logic Cells.



There are four locations along the north side and south side of each street half (west half and east half) where logic cells can be placed. Some cells require

only one cell site, others require two cell sites (see cell sizes in Table 7 at the end of Chapter 2). The clock amplifiers shown in Figure 66 require three cell sites. Most ICs do not fill every possible location with a logic cell, so the AA macro assigns a null cell (the AF cell) to all unused locations. Figure 67 reproduces that portion of the IC Boolean showing cell placement. It is this page which tells the computer Route program where to place the cells in the IC package. This page can be generated automatically by the PlaceCells program except for the placement of the output cells which must be specified by hand. Unused locations are represented by double dashes.

FIGURE 67. Cell Placement in the IC Boolean Document.

<i>Package EI Boolean Placement</i>									
Street A	(OA)	(OA)	(OC)	(OC)	(--)	(--)	(OD)	(OD)	
	(oa)	(oa)	(oc)	(oc)	(--)	(--)	(od)	(od)	
Street B	(OB)	(OB)	(--)	(--)	(--)	(--)	(OE)	(OE)	
	(ob)	(ob)	(--)	(--)	(--)	(--)	(oe)	(oe)	
Street C	(--)	(--)	(BA)	(BA)	(AA)	(AA)	(--)	(--)	
	(--)	(--)	(BB)	(BB)	(AB)	(AB)	(--)	(--)	
Street D	(--)	(ZZ)	(ZZ)	(ZZ)	(--)	(--)	(OF)	(OF)	
	(--)	(zz)	(zz)	(zz)	(--)	(--)	(of)	(of)	
Street E	(OJ)	(OJ)	(--)	(--)	(--)	(--)	(OG)	(OG)	
	(oj)	(oj)	(--)	(--)	(--)	(--)	(og)	(og)	
Street F	(--)	(--)	(--)	(--)	(QA)	(QC)	(--)	(--)	
	(--)	(--)	(TA)	(TA)	(QB)	(QD)	(--)	(--)	
Street G	(OK)	(OK)	(--)	(--)	(--)	(--)	(OH)	(OH)	
	(ok)	(ok)	(--)	(--)	(--)	(--)	(oh)	(oh)	
Street H	(OL)	(OL)	(--)	(--)	(--)	(--)	(OI)	(OI)	
	(ol)	(ol)	(--)	(--)	(--)	(--)	(oi)	(oi)	


<i>Address line drivers</i>									
FG	OA = AA + BA .								
FF	oa = aa ba .								

The Route program determines where to place the cells in the IC package by reading the boolean placement in the IC Boolean document. It determines what type of cell to place in each location by obtaining the signal type from the boolean placement and then finding what type of cell is specified for use by

that signal type from the boolean statements contained in the IC Boolean document. In Figure 67 you can see that a cell requiring two locations, identified by signal type OA, is to be placed at the west end of A street, on the north side of the street. If you look at the boolean statements reproduced in Figure 67, you can see that signal OA requires cell type FF (listed first in the boolean statement line). Thus the Route program will place an FF cell on the west end of A street, north side.

In Figure 68 the layout of the EI package after the Route program has placed the logic cells has been reproduced. However, in this illustration all of the null cells have been removed to help you compare the actual cell placement with the boolean placement information in the IC Boolean (Figure 67).

FIGURE 68. Cell Placement in EI Package without Null Cells.




By comparing Figures 67 and 68 you should be able to get a good sense for the relationship between the cell placement specified in the IC Boolean document and the resulting cell layout in the actual IC package. Notice that all of the cells

use two locations on the streets except for the clock amplifiers in west D street, which use three locations, and the group of four Qs on east F street which use one location each. Every unique IC package is different in the number and combinations of logic cells used.

In Figure 69 just four cell locations from west F street (near the avenue) have been enlarged so cell layout can be seen more clearly.


FIGURE 69. Close-Up of Cells on West F Street.



After the Route program has placed all the logic cells called for by the IC Boolean, it uses the boolean statements in that document to ascertain which signals from the different cells need to reach other cells inside the package. The program also reads the pad assignments to learn which signals need input traces from the bonding pads and which cells need output traces to the bonding pads. Once this data has been read, the Router begins to create the necessary macros which will draw the interconnecting trace lines for all of these signals. The router uses pre-determined paths in the streets and avenue in which to draw these trace lines.

A fully routed IC package can be seen in Figure 70. When the final assembly is complete, the AA macro inscribes the name of the IC in the left, bottom corner of the die.

FIGURE 70. Step Three: Routing the Final IC Assembly.




Since the resolution is poor in Figure 70, two more illustrations of portions of this fully routed IC can be seen in Figures 71 and 72. However, Figure 70 does serve to give an overall picture of a fully routed CRAY-3 IC of average complexity. Notice the small dark square below the IC in Figure 70. It represents the actual size of a CRAY-3 GaAs integrated circuit.

The trace lines which you see traveling horizontally in the streets are 2.4 microns wide. The much larger horizontal power bus traveling down the center of each street is 9.6 microns wide. Vertical trace lines drawn within the avenue are 3.2 microns wide. The large vertical bus in the center of the avenue is

actually two buses, one for the normal clock signal and one for the invert clock signal which come out of the clock amplifier outputs on D street west. These buses are each 16.0 microns wide.


In Figure 71 the resolution is such that you can finally see the individual routed lines in the streets and avenues. Figure 72 will help show them even more clearly, along with some of the individual integrated circuit components in a cell.

FIGURE 71. Bottom Left Corner of the EI Package.



Vertical trace lines are drawn on layer D (which becomes metal one in the die) and horizontal trace lines are drawn on layer F (which becomes metal two in the die). This way the routed lines are able to cross each other without shorting. When a routed line needs to change directions it goes to a different layer, using a via to connect the two different layers at a given point. The vias are actually small rectangular holes in the insulating material which separates the different metal layers. Metal is then deposited in the holes to connect the two layers. Most of the vias used in the streets and avenues are 1.6 by 2.4 microns in size.

FIGURE 72. Routed Portion of C Street West and Adjacent Avenue.



Whereas the printed circuit boards almost always require hand routing of the trace lines, the IC route programs are almost always successful in automatically routing each IC package. Only when the computer fails to completely route an IC package is manual intervention required. Generally, this involves one or two operations.

The first operation that will be tried to correct the problem is hand placement of the logic cells within the IC. Normally the router fails because either a street or a section of the avenue, or both, is too crowded and it is unable to draw the number of trace lines needed in that area. This may occur because the logic design engineer miscalculated on the number of trace lines needed for all of the cells in a street. Most of the time, however, it occurs because the router is not smart enough to figure out the right combination of connections to get all the lines into a given area. In other words, there are enough absolute paths available in the street or avenue, but because the router lacks the intelligence to use the available paths most efficiently it is unable to use all of the paths that are there and fails to route a few lines. Often this problem can be solved by moving the cells around in the IC package so as to more effectively utilize available routing space. This is done by simply editing the boolean placement page in the IC Boolean document. The IC is then re-routed using the new cell placement.

If intelligent hand placement of the cells still fails to enable the router to fully route the IC, actual hand routing of some of the trace lines will be necessary. This involves hand editing of the routing macros (BA through BH for the streets and BI for the avenue). This may involve the creation of the macros which will draw the unrouted lines and vias if the router could not place them at all. It may also involve changing the positioning of lines which the router has already placed, but placed inefficiently.

In the case where the cell placement is done by hand, the router is still used to route the IC. When trace lines need to be added by hand, the IC is first automatically routed and then the necessary macro files which the router created are edited by hand in order to either add new lines or move lines which the router has already placed, or both. Since the Route program can route an IC of average complexity in about 50 seconds on a Macintosh Quadra 700, hand routing adds a considerable amount of time to the process of completing an IC package. The time required for hand routing can take from five minutes for simple cell placement changes to a day or more for extensive routing changes.

3.3 Verification of the IC Artwork

Once an integrated circuit package has been fully routed it is ready for testing. There are four basic means of testing the completed IC:

- Tests built into the Route software program.
- A separate program (VPlat) checks all of the individual macros to be sure they are the correct overall size and that there are no internal syntax violations in the text of the macros.
- The Long Extract which checks for boolean equivalency in the routed IC and performs DRCs (design rule checks).
- Visual inspection of the IC on a color monitor or by means of printed plots.

Generally, the first three test areas are always performed while the last test (visual inspection) is usually not. Once the computer programs have been written and tested (with the testing involving visual inspection) their operation and testing routines are basically foolproof.

The Route program has several built-in tests which are automatically performed as part of the routing operation. These tests check for incorrect cell placement, illegal pad assignments, missing or unused boolean terms, routing conflicts, unassigned boolean terms in streets or the avenue, total trace line length in grid units for the trace lines between clocked cells (including length of padding, if added), and capacitive clock loading results.

VPlat is a separate program which is run using the macro data of the completed IC. It performs two basic functions. First, it knows the correct overall size for every macro in the CRAY-3 Cell Library and compares those standards with the macros in each routed IC. If it finds any discrepancies it alerts the user. It also prints an output file of macro sizes whether or not there are errors. Secondly, it checks the text of the macros for syntax errors (for example, a missing period delimiter at the end of a macro line) and alerts the user of any errors. If syntax errors are found, it prints an error message to the VPlat output file. The message includes a listing of the errors found.

The Long Extract is the most comprehensive test run on completed ICs. It gets its name from the first function it performs, that of exhaustively extracting the boolean from a completed IC. This part of the test is a type of "reverse routing" of the IC. Since the IC was created using the IC Boolean document, the Extract reads all of the macro data from the completed IC and reconstructs the IC Boolean (i.e. "extracts" the boolean) from the routed IC. It then compares its reconstructed boolean (based solely on the routed IC) with the original IC Boolean document from the logic design engineer and lists any discrepancies. An IC package is never sent out for manufacture if it fails this test.

The second test performed by the Long Extract program is a DRC (design rule check). This part of the program checks every routed trace line in one-tenth micron increments to be certain no IC design rules of spacing or proximity have been violated. It also detects dead-end traces. Because of the very high resolution of detection used in this check, this part of the program is processor intensive. Therefore, the Long Extracts for each IC package are run on a CRAY-2 or a Macintosh Quadra. On the CRAY-2 the processing time averages six minutes per IC. Those times average 47 minutes on a Macintosh Quadra 700 with 20 meg of RAM.

All of the tested data for each IC is copied to a micro-diskette which serves as the IC Master for that package type. These IC Masters are stored for reference and for temporary use during the fracturing process. The next chapter explains the process of mask production including the fracturing of a set of integrated circuit packages to make a mask set which can be used by Cray Computer Corporation's GaAs Fabrication Facility.

Integrated Circuit Mask Production

This chapter continues our discussion of the steps involved in producing the GaAs integrated circuits used in the CRAY-3. In the previous chapter we learned how the digitized data is constructed for each individual integrated circuit package. In this chapter we will follow that data further along in the process to learn how the mask data for the glass masks is prepared from the digitized data of the individual integrated circuit packages.

4.1 The Glass Plates or Masks

The mask vendors employed by Cray Computer Corporation use Perkin-Elmer E-beam machines to expose electron-sensitive glass plates. The glass plates are exposed and developed in a class one clean room. The glass plates are 12.7 cm square and have a 5X image centered on them; i.e. the image you see on the glass is five times larger than what it will be when projected onto the GaAs wafer. These glass plates are then used in the Canon steppers at the Cray Computer Corporation GaAs Fabrication Facility to expose the different layers of the digitized data onto the GaAs wafers. The glass has an anti-glare surface treatment to minimize reflections in the Canon steppers. The opaque portion of the completed plate is chromium.

There are 10 glass plates (masks) required to produce one set of ICs. The first plate provides alignment marks that will be used for aligning the subsequent plates. Alignment from layer to layer is critical when the objects being projected are measured in tenths of microns. Obviously, the registration needs

to be much more precise than that shown by some of the color photos in your morning newspaper. Plates two through nine are the various circuit layers. Plate 10 determines where a passivation coating will be placed on the finished wafer. This coating is used to protect the integrated circuits from damage after they have been created. It is omitted where the saw blade will later cut the wafer into separate die and where the gold leads will be bonded to the die bonding pads.

The E-beam machine has a variable spot size for exposing the glass plate. A spot size is chosen so it matches the address size of the data. As we learned in Chapters 2 and 3, the integrated circuits are made up entirely of rectangles on a 0.1 micron grid. Since the address size is considered to be the smallest increment that can be used in dimensioning the rectangular objects within the circuits, the address size for CRAY-3 integrated circuits is 0.1 micron at the wafer level. When enlarged five times on the glass mask, this results in a 0.5 micron address size for the data. Therefore, the E-beam spot size is 0.5 micron. The positioning of the spot on the glass mask occurs in increments of the spot size. So a 0.5 micron spot can be positioned in 0.5 micron increments across the glass.

Mask vendors use the terms dark field and light field to refer to the area on the mask that is not data; i.e., the area or field surrounding the data. For simplicity, let's assume there is one rectangle of data on a particular layer. If the rectangle is clear, then the surrounding field is dark. Thus it is called a dark field layer. However, if the rectangle is opaque (chrome), then the surrounding field is light (clear). This would be called a light field layer. Table 8 lists the various types of fields used in the CRAY-3 integrated circuit masks.

4.2 Mask Data


The data supplied by the Artwork Department for use on the E-beam machines is on half-inch magnetic tape in the MEBES format. The data is in four separate parts but is found on each glass mask. The vendor must combine the separate data as required to produce a complete mask layer. Each of the four parts is described below and illustrated in Figure 73.

The first type of data contains the alignment marks seen in each of the corners of the glass masks. These alignment marks are designed in advance by Cray Computer Corporation and supplied to the mask vendor. This data stays on file with the mask vendor so it can be added to each mask layer that needs the alignment marks.

A bar code that is automatically generated at the time the integrated circuit data is fractured comprises the second type of data seen on the glass masks. The bar

code contains 10 characters. The first five are the five digit mask number. This is followed by a dash and a four character layer title; for example, 90062-43RI. The Canon stepper scans the bar code to verify that the correct glass mask is in place in the machine.

FIGURE 73. 16-Up Mask Frame Layout.



The third type of data on a mask is the Process Control Monitor frame data. This data is also designed in advance by Cray Computer Corporation. It contains alignment patterns and PCM circuits. This frame surrounds each group of 16 IC packages. The circuits within the frame are individual CRAY-3 integrated circuit cells which can be checked electrically during the fabrication process to give an indication of problems that might arise during the processing. The PCM circuits also help the GaAs fabrication facility personnel keep the process from drifting from established process parameters. As with the alignment data, the PCM frame data usually stays the same from mask set

to mask set and is kept on file by the mask vendor. The mask vendor then combines this PCM frame data with the other mask sections to produce a complete mask layer.

The remaining section of data found on each mask layer is the IC circuit data itself. The layout arrangement currently in use clusters two sets of four-by-two ICs (two eight-ups) together to form a 16-up mask set. The 16 integrated circuits on a mask may be all the same package or any combination of different packages. Each set of four ICs has a narrow frame around it on the passivation layer. This passivation frame is automatically generated by the fracture program. The purpose of this frame is to provide a gap in the passivation material on the wafer where the wafer will be cut (diced) into individual die.


The IC circuit data and the bar code data are unique to each mask set. Therefore, these two sets of data need to be sent to the mask vendor each time a mask set is ordered. The alignment marks, PCM frame and passivation information is not unique to each mask set and can generally be used for all of the IC mask sets used to make the ICs for the CRAY-3. The exceptions are when any revisions are needed to this information or when special test ICs are made which require special probe pads. In the latter case, the passivation information must be changed to allow for these special pads in addition to the normal 52 bonding pads around the periphery of each IC package.

4.3 Fracturing

The data for each IC that will make up a particular mask set is loaded from the IC Master diskettes. The Fracture program allows the user to place the individual ICs in the proper order on the mask. Once all the data is properly loaded from the diskettes and the mask number specified, the fracturing process begins. The IC data is fractured to conform to the MEBES standard. Basically, fracturing means that the IC data represented by rectangles of various sizes is broken down into smaller data blocks; i.e., the data is fractured, or broken, into blocks of a specific size. However, before the Fracture program fractures the IC data it performs data rotation.

The data rotation is necessary to properly align the FETs of each integrated circuit with the GaAs crystalline structure and to provide the proper imaging for the Canon steppers (similar to inserting a 35 mm slide in your projector so that the images are not upside down or facing the wrong way when projected on a screen). This data rotation also places the bar code where there are no grabber fingers in the mask vendor's processing.


FIGURE 74. Data Rotation During Fracturing.



All of these factors added together results in a single 180° rotation about the Y-axis prior to the data being fractured. This is easily accomplished with the Fracture program. The program has a default setting that requests a 180° rotation about the Y-axis.

The one factor that warrants further explanation is the orientation of the FETs in relationship to the crystalline structure of the wafer. The wafers have a flat side which indicates a polarity of the lattice structure within the GaAs wafer. The desired orientation of a FET is to have the width of the gate (which is the longer dimension) perpendicular to the flat of the wafer. This is illustrated in Figure 75.

FIGURE 75. Orientation of the FET Gates on a Wafer.




Once the IC data is properly rotated it is fractured to the MEBES format by the Fracture program. The area of data to be fractured is broken down into narrow horizontal rows called stripes. This area to be fractured is further broken down into wide vertical columns called segments. Therefore, one block of data is the height of one stripe and the width of one segment. One block of data contains whatever portion of the rectangular shapes comprising the IC that happen to reside in that area. The fracturing of a 16-up mask set is illustrated in Figure 76.

FIGURE 76. Fractured Stripes and Segments.

16-Up Mask Set

Note: In this illustration the 16 die are all the same. However, they can be any combination of the same or different die.



The stripes are 512 address units high and 32K address units wide. Address units are 0.1 microns. This divides the total area to be fractured into 312 stripes by five segments.

After the IC data is fractured, it is written in the MEBES format on half-inch tape at 1,600 BPI, phase encoded. The tape is magnetically un-labeled. A typical tape containing fractured IC data has 17 files. The first file is an Archive Header File. It contains the names of the subsequent data files and their length. It also has a tape ID ("CRAYIC") and the date the tape was generated.

Following the Archive Header File, there is one data file for each layer which was fractured. Typically, there would be eight layers and thus eight data files for a new mask set. This includes the standard seven data layers and the automatically produced proton isolation layer. Each of these eight files contains a header record with the file name, date created and fracturing parameters (such as address and spot size, stripe height and fracturing window size). The layer data in MEBES format would follow the header record for each layer.

Finally, the tape would contain the bar code files. Generally, there is one file for each fractured layer. Therefore, a standard tape would have eight bar code files.

The 16 data files, which we have just discussed, have a standard naming convention. All file names are 12 characters long. A nine character identification precedes a period and a two digit layer number. This is illustrated in Figure 77.

FIGURE 77. Fractured Data File Name Conventions.

IC Circuit Data Files	Bar Code Files
<u>CRAY 90062 . 43</u> └── 2 digit layer number └── Mask Number └── File type (CRAY = IC data)	<u>BARC 90062 . 43</u> └── 2 digit layer number └── Mask Number └── File type (BARC = bar code)

4.4 Mask Layers


There are 10 glass masks required to process a standard set of ICs. As we have seen, two of these glass masks (alignment and passivation) do not change from one IC set to the next, except in special cases. The other eight glass masks are unique to each IC set that is made. All 10 mask layers are described in Table 8. They are listed in the order in which they are used in the IC fabrication process.

TABLE 8. Mask Layer Data.

Cell Layer	CrayDraw Layer	Circuit Data	GaAs Fab Layer	Frame Data	Layer Title	Description
N/A	N/A	N/A	1	Clear	1A-A0	Alignment
A	1	Clear	2	Clear	20ND-A0	DFET N- Implant
B	2	Clear	3	Clear	3N+-A0	N+ Implant
C	3	Clear	4	Clear	4OC-A0	Ohmic Contact
N/A	9	Chrome	41	Chrome	41I-A0	Proton Isolation*
H	8	Clear	43	Clear	43RI-A0	Recess Insulator
D	4	Clear	5	Clear	51SM-A0	First Metal (Schottky)
E	5	Clear	6	Clear	601V-A0	Vias
F	6	Clear	8	Clear	652M-A0	Second Metal
G	7	Clear	8	Clear	8S-A0	Passivation
*The proton isolation layer is automatically generated in the fracturing process. The N- and N+ layers (CrayDraw layers 1 and 2) are combined and then oversized by 0.8 μ m. The layer does not exist in CrayDraw until it is created during fracturing.						

The first mask in a set (the alignment layer) contains only alignment marks. It does not contain PCM frame data, IC circuit data, or a bar code. The last layer (passivation) contains PCM frame data and IC circuit data, but no bar code. The other eight layers that are unique to each mask set have PCM frame data, IC circuit data, and a bar code. The descriptions “clear” or “chrome” in Table 8 are referring to the data itself and not to light and dark fields. Figure 78 illustrates how the layer title is read.

FIGURE 78. Understanding the Layer Title.



The last two characters of the layer title warrant further explanation. The first of these two characters allows for more than one set of glass masks for the same mask number. This might be done if there is a high volume of wafers needed for a particular IC mask set. Thus, it might be advantageous to have more than one mask set on hand so that both could be in use at the same time. Another reason to have more than one set of glass for the same IC mask set is to have redundancy in the event one set is broken.


The last of these two characters allows for a modification to a layer within a mask set. For example, suppose an error was discovered that only affected the second metal layer. Rather than making all layers of glass again, it would be less expensive to make just this modified layer. Then the layer revision number would be changed to identify this as the modified piece of glass.

4.5 Positioning of ICs from Mask to Wafer.


Because of the data interpolations involved in the fracturing process, the correlation between the position of an individual integrated circuit package in the fractured mask data and the resulting position on a wafer can be quite confusing. In this brief section we will try to bring some clarity to this question of integrated circuit positioning.

When the IC data is first loaded into CrayDraw, in preparation for fracturing, the resultant positioning of the integrated circuit packages on the computer monitor is a mirror image of the positioning they experience on the wafer if you have the flat of the wafer at the bottom and view the die. After the IC data is fractured, the ICs again appear on the computer screen as a mirror image of their positioning on the wafer with the flat of the wafer at the bottom; in other words, flipped from right to left compared to how they would appear on the wafer. This is illustrated in Figure 79.


FIGURE 79. Correlation of IC Positioning.



IC Packages as they appear on the computer screen after being entered into CrayDraw for fracturing.



IC Packages as they appear on the computer screen after reading them back into CrayDraw after fracturing.



The same IC Packages as they appear on the finished wafer.



Black rectangle indicates position of Pad One and IC Name.

The four areas on the wafer where the die are not lettered are reserved for drop-ins.

The only difference between the views of the ICs on the computer monitor before and after fracturing is the addition of the fracture stripe and segment boundary lines in the second view (after fracturing). Other than that, the IC data itself is the same, and the orientation is identical. The positioning will be mirrored, as will the name inscriptions. In other words, the letters of the inscriptions will be backwards.

The field numbers on a wafer can be determined by looking at Figure 79. The fields are designated by combining the numerical digits of the Row with the numerical digits of the Column. For example, the first active field at the bottom left of the wafer would be field 105 (or 10,5). That is, the field rectangle containing die types A, B, C, D, E, F, G and H lies in Column 10 and Row 5. Therefore, it is in field 105. Since the ICs are always loaded into CrayDraw in alphabetical order, one can determine which die of any mask set are located in whatever field by using Figure 79. This is true—as illustrated—for 16-up mask sets. Eight-up mask sets would just repeat the first eight die types for the second group of eight; i.e., the group of die lettered I through P would also be the A through H die again.

In the next chapter we follow the glass masks to Cray Computer Corporation's GaAs Fabrication Facility to get a basic understanding of the processes involved in fabricating a CRAY-3 GaAs integrated circuit.

Integrated Circuit Fabrication

Up to this point, we have presented material on the design of the individual components that make up an integrated circuit package. We have also examined how the data for producing an IC is created, tested and fractured so that reticles (glass masks) of all the layers of data that make up an IC can be made. In this chapter we will follow the use of these reticles in the actual fabrication of the integrated circuit GaAs wafer in Cray Computer Corporation's GaAs Fabrication Facility.

5.1 Integrated Circuit Fabrication

The integrated circuits used for logic packages in the CRAY-3 computer system are depletion mode, gallium arsenide Metal Semiconductor Field Effect Transistors, or MESFETS. The gallium arsenide refers to the material used for the integrated circuit substrate. The field effect transistor is the primary switching device used in the CRAY-3 logic design. A second device, a Schottky barrier diode is also used extensively. Both devices were discussed in Chapter 2.

The gallium arsenide material used as the substrate for the integrated circuits is grown as a ten centimeter diameter, single crystal in the form of an ingot. The material is then sawn into 635 micron thick wafers and polished. A flat is ground on one side of the wafers in order to orient the wafer correctly during processing. This is necessary because GaAs has crystallographic orientation properties which give rise to different electrical operating characteristics in

different directions. More specifically, transistors must be oriented so that the width of the FET is perpendicular to the flat of the wafer (see the discussion in Chapter 4 and Figure 75). This orientation gives better and more consistent device performance when used in conjunction with the current Cray Computer Corporation's GaAs fabrication process.

Transistors and diodes are made on the surface of the GaAs wafer by selectively implanting ions with a linear accelerator to provide locally conductive regions for FETs, diodes, capacitors and resistors. The resistivity of the GaAs without an implant is very high—it is essentially a non-conductive insulator. The un-implanted regions (called “semi-insulating”) between the implanted active devices provide the necessary electrical isolation for the integrated circuits. Metal layers are then added to the surface for contacts and interconnections on and between the circuit components. Two levels of gold metal interconnect are used in the CRAY-3 integrated circuits.

Patterns of implanted areas and patterns of contacts and metal layers are processed on the GaAs wafer by coating the wafer with a photo sensitive material called “photoresist” and then exposing the photoresist with selectively masked light patterns. The photoresist is chemically removed (developed) where the light has chemically altered the photoresist material. The remaining photoresist material is then used as a mask for ion implanting and etching of the various dielectric areas, metal contacts and interconnects comprising the micro-circuits.

Ions are implanted in the wafer by accelerating the ions in a vacuum at very high velocities in a strong (80-170 keV) electrostatic field and essentially embedding them into the wafer surface. The implant is done in a selective manner by using a photoresist mask material which will stop the ions before they reach the wafer surface. The openings in the photoresist define the locally implanted areas which will become the various devices after fabrication.

Glass layers (dielectrics) are applied to the wafer surface by sputtering and plasma enhanced chemical vapor deposition or PECVD. The sputtering process involves moving atoms of the source (target) material to the substrate (wafer) material in a moderate vacuum. The target and substrate are usually spaced a few inches apart. The atoms of the source material are freed from the surface by bombarding the source with energetic ions provided by a radio frequency induced plasma. The ions can either be inert (e.g. Argon) or reactive (e.g. Nitrogen). In the case of Argon, the process of sputtering atoms from the target with energetic Ar ions and transferring them through a moderate vacuum to a substrate is largely mechanical in nature. No actual chemical activity takes place. In the case of sputtering with a reactive gas like Nitrogen, the N_2 both sputters (knocks) atoms off the target (e.g. Silicon) and then also chemically reacts with the Silicon in the plasma or on the substrate material. Specifically, in the Cray fabrication process a Silicon Nitride (Si_3N_4) dielectric is deposited using a “reactive sputtering” process incorporating a N_2 plasma with a Silicon

target. The travel of the individual atoms from source to target is sufficiently random that material is applied to the wafer surface from many angles. This makes for generally good coverage of surface steps on the wafer.

The second method used to deposit dielectrics in the Cray GaAs fabrication process is the use of plasma enhanced chemical vapor deposition, or PECVD, to deposit Silicon Dioxide (SiO_2). This process is also conducted in a moderate vacuum utilizing a plasma. In this process the source materials are not provided by a solid target, but are provided by gases. The plasma enhanced, chemically deposited SiO_2 uses Silane (SiH_4) and Nitrous oxide (N_2O) as its source gases. The SiH_4 and N_2O chemically react to form SiO_2 in the plasma at approximately 250°C . The SiO_2 is subsequently deposited on the substrate (wafers). The glass or dielectric layers deposited on the GaAs wafers provide a solid medium in which to fabricate the micro-circuits. The insulating properties of the Si_3N_4 and SiO_2 dielectrics also provide ideal electrical isolation between devices and in between the two levels of interconnects.

Some metal layers are applied to the wafer with sputtering and some with electron beam evaporation. Electron beam evaporation involves moving the molecules of the source material to the target material in a high vacuum. The source and the substrate (wafer) are spaced about 50 millimeters apart. The source is bombarded with electrons which heat the source material to its melting point. The individual molecules of the source material then evaporate and travel through the vacuum and are deposited (condensed) on the wafers. The travel of the individual molecules from source to target in this case is a relatively straight line. Most molecules arrive at the wafer surface at right angles and step coverage is usually poor.

The primary metal used in the integrated circuit process is gold. Gold is chosen because it is an excellent conductor of electrical energy and it is at the same time chemically inert. In some cases, thin layers of Titanium (Ti) and Platinum (Pt) are also deposited before the gold in order to provide high quality Schottky barriers on the GaAs surface. A Gold Germanium (AuGe) alloy is also used in small quantities to form low resistance ohmic contacts to GaAs.


Some metal steps in the fabrication process use metal lift-off. This involves placing the patterned photoresist on the wafer surface first, and then covering the surface with the metal layer. That portion of the metal which is over photoresist is then removed by chemically dissolving the photoresist. The remaining metal forms the various contacts and interconnects used to fabricate the micro-circuits.

The approach and methods described here are integrated into a sophisticated process sequence comprised of 10 masking levels to fabricate the GaAs integrated circuits making up the CRAY-3 logic packages. These process steps will be discussed in more detail after we introduce the Canon wafer steppers.

5.1.1 Canon Wafer Stepper

The machine which does the photolithography in the integrated circuit fabrication process is a Canon 5X wafer stepper. This machine uses glass plates called reticles with chrome patterns which are five times the linear dimensions of the desired circuit patterns on the GaAs wafers. Ultraviolet light is then used to project the image of the circuit pattern through a 5X reduction lens onto the wafer. The reticles used in CRAY-3 circuit work are 12.7 centimeters square.


FIGURE 80. Stepping Sequence on a GaAs Wafer.



The Canon photolithography machine is called a wafer stepper because the pattern projected on the wafer covers only a small portion of the available wafer surface (see Figures 80 and 81). The wafer is then moved, or stepped, and the pattern is repeated. Each reticle may contain the layout for 4, 8 or 16

different circuit types. The Canon wafer stepper is capable of allowing up to 16 different circuit types, if required.

FIGURE 81. Stepped Die Patterns on a Wafer.



The various reticles, containing the chrome circuit patterns, are mounted in cassettes located within the Canon projection system. The reticles are automatically pre-aligned in the Canon 5X optical path using reticle pre-alignment marks located on the outer edges of the plates prior to the precise wafer-to-reticle alignment.

The Canon wafer stepper has numerous optical and electronic feedback systems which automatically align and focus projected circuit patterns to previously imaged patterns on the wafer. This system involves both alignment marks on the wafer and alignment windows on the reticles. The alignment marks are fabricated with Titanium (Ti) placed on the wafer in each stepped

field (see Figures 80 and 81). These marks are placed on the wafer in the first process step which is exposed in a blind (non-aligned) step and repeat process. This is accomplished with a high degree of accuracy using the Canon wafer stepper laser interferometer stage. Subsequent steps in the process involve exposure and alignment using a feedback system for each of the fields on the wafer. The alignment process involves projecting a red light through the alignment windows in the mask. The wave length of this light does not expose the photoresist material. It therefore does not alter the permanent alignment marks on the wafer. The alignment windows are positioned in the glass mask so that the red light strikes the alignment marks and is reflected back to photosensors in the servomechanism system. The wafer chuck position is then corrected slightly to align the image perfectly with the previous patterns on the wafer. The ultraviolet light is then turned on to expose the photoresist for the current process step. This positioning system means that each of the individual Canon fields has its own alignment tolerances.


5.1.2 GaAs Fab Processing Steps

The gallium arsenide processing steps discussed in the next paragraphs are illustrated in Figures 82 through 87. The illustrations should be studied often as the text is read to gain a better understanding of the processes involved.

The process of making a GaAs integrated circuit begins with a Silicon Nitride "cap" (see Figure 82). This is essentially a thin layer of glass which is used to protect the surface during processing. The glass is sputtered onto the wafer. This step is done independently of the circuit specific process steps and is generally not included in a process step list.

The first step involving photolithography is the creation of alignment marks. The alignment marks located on each side of the circuit patterns are used by the Canon stepper to align later mask levels in the process. These marks are not circuit specific and the same glass mask is used for all circuit types. The marks are an array of alternating highly reflective (Titanium metal) and low-reflective (GaAs) patterns.

FIGURE 82. Dielectric Encapsulation and N- Implant.




GaAs Substrate is 635 microns thick.
Illustrations are not to scale.
Vertical scale is greatly exaggerated.

The second step in the process is a light (low dose) ion implant illustrated in Figure 82. This is called the N- implant. This is the implant which makes the semiconducting channels for the transistors. The implant is called “light” because both the energy and the dose of the ions is low allowing the ions to remain near the surface of the wafer. This extremely thin semiconducting channel is optimized for very high-speed operation. Resistivity of this implant is about 460 ohms per square. The ions used are singly ionized Silicon (Si^+).

The third step in the process is a heavy ion implant illustrated in Figure 83. This is also called the N+ implant. This implant is used to make contacts at the ends of the transistor channels and is also used in diode structures. Ions used are doubly ionized silicon (Si^{++}). In contrast to the light channel implant, here the energy is higher and the dose is greater. These ions are driven deeper into the wafer surface. Resistivity of this implant is about 365 ohms per square.


FIGURE 83. Heavy Ion (N+) Implant.



At this point in the process the wafer is rapidly, thermally annealed (RTA). The temperature is raised to about 850° C for approximately two minutes. This anneal allows the silicon atoms to become substitutional in the GaAs crystal structure, thereby allowing the GaAs to become electrically conductive. This process step requires the use of the Silicon Nitride cap in order to protect the GaAs surface during the high temperature implant anneal. If the GaAs surface was not sealed, arsenic would diffuse or vaporize above temperatures of 600° C.

The fourth step in the process is the addition of ohmic contact metal shown in Figure 84. This metal is intended to make a good ohmic contact (low resistance) with the N+ implant in the wafer. In other words, this is the metal connection to the end of an implant channel. This metal is a mixture of gold, germanium and nickel. This metal mixture is deposited with electron beam evaporation onto a photoresist patterned wafer. The metal which is over the photoresist is removed by the lift-off process. The wafer temperature is then raised to about 450° C for 30 seconds. This temperature causes the newly added metal mixture to melt and form an alloy, wetting the doped GaAs surface making a very low resistive ohmic contact.

FIGURE 84. Ohmic Contact Metallization and Proton Isolation.




The fifth step in the process is called proton isolation. This is an implant of Hydrogen ions (H⁺). This implant is used to damage the GaAs crystal structure in areas where current flow is not wanted. The GaAs substrate already has a high resistance, but the resistivity from ingot to ingot can vary considerably. Proton isolation is used in order to more predictably provide high resistivity regions between devices ($\geq 10^6$ ohms per cm). The damaged regions also provide high numbers of electron traps which help suppress sidegating or backgating leakage. Leakage is an undesirable feature in FETs. The voltage potential on one device may influence the electrical behavior of a closely

spaced adjacent device. Proton isolation minimizes these undesirable effects of leakage.


The sixth step in the process is called first layer metal or Schottky metal (SM) and is illustrated in Figure 85. It is referred to as “Schottky metal” because this metal, when located over an implanted region creates a Schottky barrier diode. It is also called gate metal because it forms the gate in a transistor (FET). This metallization is deposited with electron beam evaporation. A sequence of metals (titanium, platinum and gold) are applied. Unwanted metal is removed by a lift-off process as in the case of ohmic metal.

FIGURE 85. Schottky Barrier and First Layer Metal Interconnect.




The seventh step in the integrated circuit process is the addition of inter-level dielectric glass. This is a silicon dioxide glass and is needed to separate the first and second level metal layers electrically. This material is deposited by PECVD onto the wafer. Selected areas of the glass are then plasma etched to provide openings where first-to-second layer metal connections are desired. These areas are called vias. The vias are then filled with gold by electron beam evaporation to return the wafer surface to a near planar condition. The unwanted metal is again removed by the lift-off process. The same photoresist serves for the plasma etch step and the metal lift-off step of this process sequence.

FIGURE 86. Second Level Metal.



The eighth step in the process is the addition of second level metal (see Figure 86). This metal is also electron beam evaporated onto the wafer. The metal which is used consists of titanium and then gold. This metal layer is thick and requires a special retro-grade photoresist profile for the subsequent metal lift-off process. Again, the unwanted metal is removed by the lift-off process.

FIGURE 87. CERMET (Optional) and Passivation.



The last step in the process is the addition of silicon dioxide glass to protect or passivate the surface of the completed integrated circuit. This glass is deposited onto the wafer by the PECVD process. Selected areas of the glass are then removed by plasma etching where bonding pad connections are desired. This step is not circuit specific because the openings in the passivation glass for bonding pads are the same for all GaAs circuit packages used in the CRAY-3, except for special test ICs. Thus, this mask is re-used for all circuit types.

TABLE 9. GaAs Fab Processing Steps.

Step	Mask Used	Name	Description
Preliminary	None	Deposition of a Silicon Nitride Cap	Deposits a thin layer of glass on the wafer to protect the surface.
Step One	1A-A0	Creation of Alignment Marks	Used by the Canon steppers to align later mask layers.
Step Two	20ND-A0	N- Implant	Makes the semiconducting channels for the transistors.
Step Three	3N+-A0	N+ Implant	Makes the contacts at the ends of the transistor channels.
Intervening	None	RTA	The wafer is rapidly thermally annealed allowing the GaAs to become electrically conductive.
Step Four	40C-A0	Ohmic Contact Metal	Allows a low resistance contact with the N+ implant.
Step Five	41I-A0	Proton Isolation	Helps damage the GaAs crystal structure where current flow is not wanted.
Step Six	51SM-A0	First Layer Metal (Schottky)	Deposits a metal trace layer for signal routing.
Step Seven	None	Inter-Level Glass	Electrically insulates layer one metal from layer two metal.
Step Eight	652M-A0	Second Layer Metal	Deposits a second trace layer for signal routing.
Step Nine	8S-A0	Passivation	Protects the surface of the integrated circuits except where electrical contact is needed.

5.2 Wafer Processing

Completed wafers are sent on for processing in the Cray Computer Corporation's production and manufacturing facility. Here the individual integrated circuits are tested while still in wafer form. Then the wafers are ground to a specified thickness and cut (diced) into the individual die. Once cut, the individual die are automatically placed into small waffle packs based on the test information generated earlier. But before the die can be mounted on the CRAY-3 printed circuit boards, they must have tiny wire leads attached to them which enables them to make mechanical and electrical contact with the printed circuit boards used in the machine. This process is done in Die Bonding after which the die are placed in Inventory to await their use in Module Assembly.

5.2.1 FAST Testing

The manufacturing principle for the CRAY-3 requires that the process be started with components that are known to be functional. The printed circuit boards are tested point-to-point and on a "bed of nails" tester to ensure 100 percent electrical functionality. The integrated circuits are also tested by function at full-speed in a specially developed process called Functional At-Speed Testing.

FAST is a process whereby logic circuits, clock circuits and memory circuits are tested for functionality at-speed; i.e., at 500 mHz and at speeds slower and faster to test for speed tolerances. Every input and every output of every circuit is exercised while the circuit is still part of an incoming wafer. The three major types of circuits are sufficiently different that each requires a special set of test equipment. The description that follows pertains principally to gallium arsenide logic circuits.

Wafers are tested through the FAST system as the first step in the wafer processing sequence after their arrival from the respective foundries. A full wafer record from the fabrication facility is kept in a folder in the production clean room. This record is updated with FAST results as they are made available.

The tester setup is an Electroglas wafer probe station equipped with a thermally controlled chuck. The wafer is moved around by the servo table it resides on so that each integrated circuit's 52 bonding pads can be aligned directly under 52 probe tips which move down to make contact with the pads.

Each circuit type requires a separate program for testing. This program governs the input and output acquisition, and the sequence and timing of signals. The programs are written based on the original IC Boolean design for each integrated circuit (see Chapter 3). The electronic signal generation and

acquisition for the FAST tester is controlled by a Macintosh computer which contains the programs for testing the individual circuits. The Macintosh also gathers and stores the results of each test.

In a development mode, the test of each circuit can show a variety of performance parameters for the circuit. During normal production, the tester shows good circuits, bad circuits and high current or shorted circuits. Once a wafer has been tested, the Macintosh program will develop a map of the wafer showing the results for each die in each field position. This map is later used by the pick and place machine to properly sort the individual die into the right categories, placing like categories in the same wafer packs.

Alignment of the probe tips on each individual circuit is critical. Since the Canon stepper has a high degree of accuracy, the alignment of the probe tips to the die needs to be performed only once per wafer. Obviously, incorrect alignment would give unorthodox test results. Therefore, operator attention is necessary throughout the testing process to ensure that alignment is maintained. The alignment of the die under test can be checked through a microscope on the Electroglas prober station.

Clock distribution circuits (the ZA die) are tested on a separate test station developed specifically for this purpose. The method of operation is similar to that described above. Clock circuits are normally occupying the whole of a wafer rather than being mixed with circuits of other types.

Memory circuits for Common Memory and local memory are tested on memory test stations. This test station establishes functionality and performance levels for all addresses in the individual circuit's memory.

5.2.2 Grinding

Wafers from the GaAs fabrication facility process are too thick to be used in the module assemblies. Therefore, before the wafers can be cut into individual die they must be ground from 0.635 mm to 0.208 mm in thickness.

The grinding machine is capable of multiple and individual wafer feeding. Dummy or mechanical wafers are used for calibrating the machine before submitting functional wafers for processing.

The process of grinding begins with the cleaning of the wafers prior to the calibration of the grinding machine with mechanical wafers. The wafers are mounted, face down, on a ceramic chuck. The chuck is porous which allows a vacuum to hold the wafer securely in position. After cleaning and calibration, the wafer to be ground is mounted on the ceramic chuck. Three grinding wheels successively remove surplus gallium arsenide from the back of the wafer until the desired thickness is obtained. Gallium arsenide particulates are

removed by a constant supply of de-ionized water which floods the grinding area.

Upon removal from the grinding machine, the wafers are once again cleaned. The thickness and evenness of the wafer is checked. At this point there will be very slight, circular marks visible on the back of the wafer from the grinding machine.

5.2.3 Dicing

The thinned wafers from the grinding machines are diced into individual die on a dicing machine. The dicing machine is a small circular saw which has material presented to it in a controlled fashion on an x-y table. The saw blade is mounted under a television camera having a microscopic lens. The sawing area is flooded with de-ionized water during the sawing process to take away particulates. An operator is able to watch the operation and adjust the saw cut by means of a television screen display.


The saw blade is two inches in diameter and rotates at 33,400 revolutions per minute. The blades are tapered such that the initial cut width is 15 microns and varies up to 20 microns. The blades are made of resin and are coated with two to four micron diamond grit. New blades are dressed by cutting scrap silicon wafers prior to use on production material. Once a blade is dressed, it is good for cutting a varying number of wafers depending on the constitution of that particular blade. The dicing machine automatically adjusts the depth of the cut as the blade wears.

The wafer is mounted on an aluminum frame with a pressure sensitive Mylar sheet bed. The aluminum frame has a 15.24 cm square exterior and a 12.7 cm round interior. The Mylar sheet is 0.127 mm thick and is bonded to the back of the aluminum frame. The wafer is bonded to the Mylar bed with the face of the wafer up. This carrier is used to hold the wafer during the sawing operation. A similar process is used for gallium arsenide and silicon wafers.

The saw blade is lowered from above the wafer and the current carrier frame. The cut is made in a single pass. The cuts are made with the blade extending a few microns into the Mylar tape. This is to allow complete separation of the individual circuits. The first cuts are made parallel to the wafer flat. This is the direction in which gallium arsenide tends to chip the most. The wafer is presented to the blade at a cutting speed of 5 mm per second. The second cuts are made perpendicular to the wafer flat. The cutting speed again is 5 mm per second.

The arrangement of die on a wafer is such that the spacing between the stepped groups requires two dicing cuts. Within the group or steps the die are separated by narrower streets which require a single cut as is illustrated in Figure 88.

FIGURE 88. Wafer Dicing.



There are 344 circuit package die resulting from the dicing of one ten centimeter (about four inches) GaAs wafer. Each circuit die is 3.835 by 3.835 mm, less the width of the saw cut. The saw cut is arranged such that there will be no chipping damage beyond the ohmic metal fence (crack stopper) around the functional circuits of the die. The die are 0.208 mm thick. The PCM circuit fields and alignment fields on the wafer are discarded. The individual die on the Mylar tape are then cleaned with high pressure deionized water to remove any remaining particulates left from the dicing process.

5.2.4 Pick and Place

After the wafers have been cut into the individual die they are still in the wafer format but held together by the adhesive tape which was applied just before dicing. During pick and place these segmented wafers have the individual die picked off the adhesive backing and placed into small plastic boxes called "waffle packs." Each waffle pack can hold up to 64 individual die in their own small, square compartment. The wafer is picked automatically by a machine

guided by the test data obtained at FAST. From this data the machine knows the location and type of all the good and bad die tested for that wafer. The die which have been automatically picked are also placed in the waffle packs automatically according to type and test results. The waffle picks are then delivered to Die Bonding to have the small gold leads bonded to the 52 bonding pads located around the periphery of each die.

5.2.5 Die Bonding


The die from the wafer dicing operation are next processed for the addition of gold leads. The leads are made from soft gold wire which is 0.076 mm in diameter. There are 52 bonding pads on each die which have a sputtered soft gold finish. The objective of the die bonding operation is to form a gold-to-gold bond between the wire and the pad. The technique used to make this bond is known as thermosonic ball bonding. This technique welds the gold together using heat, force, time and ultrasonic vibration. The machine which does this operation is called a Hughes Automatic Ball Bonding machine.

The bonding machine has a thermally controlled x-y positioning bed which is used to position the die for bonding. The die is loaded on the bed, which is set to 170° C for GaAs die and 300° C for silicon die, in a fixture which holds the single die by mechanical clamps and vacuum. The Hughes bonding machine is equipped with a vision system which can recognize the die patterns without human intervention and position each die for processing. The bonders also have the ability to reference each die accurately regardless of rotation of the die.

The gold wire used in the bonding process is fed from a supply spool through a nitrogen filled tube and exits through a ceramic capillary. The bonding process involves forming a small gold ball, slightly larger than the wire diameter, at the end of the wire. This is accomplished by creating an electrical arc between an electrode and the wire tip. A high temperature plasma, about 5000° C, is created during the electrical discharge for a controlled length of time. This forms a precisely sized ball for the next step.

Following ball formation, the ball is lowered to the gold bond pad with a controlled force. Once the specified force has been reached, the ceramic capillary is energized with an ultrasonic vibrating motion which, together with the elevated fixture temperature of 170° C, welds the gold ball to the bond pad surface. The combination of force and ultrasonic energy during the welding process deforms the gold ball to an approximate size of 0.0254 mm in height and 0.152 mm in diameter.

FIGURE 89. Notcher Mechanism.



The capillary is then withdrawn from the die surface exposing the gold wire as the head is raised. At this point an automatic notching mechanism moves into the area of the exposed gold wire and strikes both sides of the wire with a steel blade. This is essentially a scissors action which cuts most of the way through the gold wire. The notch is made 0.432 mm above the surface of the die. This step in the process is illustrated in Figure 90.

A clamp then closes on the gold wire above the capillary and the head is withdrawn until the gold wire breaks at the notched point. This serves several useful purposes. The gold wire is straightened by the stretching force and stands perpendicular to the die surface and the bond is tested for adhesion to the bonding pad. The lead is terminated at a height of 0.432 mm above the die surface. Current bonding rate is 45 seconds per die.

FIGURE 90. Die Bonding Illustration.





FIGURE 91. Gold Leads Bonded to a Portion of a Die.



After the die are bonded with all of the 52 gold leads they are inspected and returned to Inventory in their labelled waffle packs for later use at Die Attach.

In the next chapter we begin to take a somewhat broader view of how these very small parts fit together with relatively larger parts to make up an

electronic module for the CRAY-3. As we take this step back to gain a broader perspective, one needs to remember that the logical operations of the machine are at heart performed within these very small gallium arsenide die which are now ready to be attached to printed circuit boards. We will pick up that step in Chapter 7.

Module Structure


In Chapter 6 our discussion turns to the assemblage of components that make up a CRAY-3 module. First an overview is presented to give a better perspective of where and how the individual components fit into the module assembly. Then, following a brief discussion of how the printed circuit boards are manufactured, the smaller parts which are part of the module assembly will be looked at individually in much greater detail.

6.1 Overview

A module assembly is the basic unit in the electronic assembly section located at the top of the CRAY-3 System Cabinet. Each module is identical in its overall size and basic mechanical organization. There are ten types of module assemblies in the CRAY-3. The modules differ from type to type in the population of integrated circuits that they contain and in their associated interconnections in the circuit boards, logic plates, power plates and resistor plate.

The module assembly is 120.9 mm wide (4.76 inches), 107.2 mm high (4.22 inches) and 7.14 mm deep (0.28 inches). A parts explosion of a module assembly is illustrated in Figure 92. The left edge of the module consists of four power blades. These machined metal blades are both the mechanical connection to the octant frame and the electrical connection to the power supply buses.

FIGURE 92. Exploded View of a CRAY-3 Module.




The center of a module assembly is made up of three plates of full module size. The center of the module assembly and the center plate of the three plates contains the terminating resistors or Ohmega layer. On either side of the terminating resistor plate are two power plates. These plates deliver the electrical current from the power blades to the circuit board stacks. These three plates extend slightly beyond the other circuit boards of the module assembly on just the left side of the module to give surface area for the attachment of the power blades.

Above and below the power plates are the two logic plates. These plates provide logic signal communication between the board stacks. Above and

below the logic plates are the two layers of sixteen circuit boards which form the stacks. These boards contain integrated circuit packages and interconnections between the packages. Electrical communication between individual boards and board stacks is via the logic plates.

FIGURE 93. Cross Sectional Diagram of a Module.




The outer surface of the module assembly is covered by a protective overlay. The overlay is arranged as part of each board stack on the top and bottom. It is made up of 2.54 cm square Ultem pieces. This entire sandwich of components is illustrated in Figure 93. It should be noted that the sandwich of layers shown in Figure 93 shows only the basic layer groupings. Each of the layers shown also has more layers of their own. For example, the logic plates themselves contain 15 layers of material each.

Logic communication and power distribution in the x-y dimension is within the circuit boards and within the plates. Communication through stacks in the z axis is via twist pin jumpers which provide the connection between logic boards, logic plates and the terminating resistor plate. Z axis connection for power is provided by a larger twist pin for power purposes which provides connection between the circuit boards and the two power plates. Both the logic and power pin jumpers make connections by means of spring forces. The assembly pins have a counter-sunk head on one end and make an interference fit at the other end in the protective overlays.

The small squares in the module assembly represent the integrated circuit packages. There are sixteen GaAs integrated circuit packages on each circuit board. With sixty-four boards in a module assembly this means that there are 1,024 positions in a module assembly for GaAs integrated circuit packages.

FIGURE 94. Arrangement of Stacks and ICs.



The integrated circuit packages are attached to the circuit boards by means of 52 gold leads which are bonded to the 52 pads on each integrated circuit (see


Figures 90 and 91 in Chapter 5). A vertical lead from each pad is produced whose length is slightly longer than the thickness of the circuit board. The attachment process is made by compressing the lead inside the gold-plated copper barrel of the printed circuit board jumper hole so that the soft gold lead is compressed to approximately two thirds of its original length, plugging the barrel with a solid plug. This forms a gold-to-gold connection. The 0.076 mm gold leads are 0.483 mm long. They are compressed inside a barrel 0.381 mm by 0.127 mm. Fifty-two connections of this type are made for each 3.835 mm square die. These 52 connections form the mechanical and electrical connection of the die to the circuit boards.

Three types of plastic spacers are used in the modules. Center spacers are used on logic boards to maintain a 300 micron gap between the top of the die and the adjacent circuit board or logic plate. These center spacers also prevent integrated circuit package damage from board warpage during assembly. The cool Fluorinert passes through the module assembly principally in the 300 micron gaps above the integrated circuit packages. The cool Fluorinert in these channels bathes the back of the integrated circuit packages as well as the gold jumpers. These two areas of contact with the fluid are the major avenues of heat transfer. Corner spacers serve the same purpose as center spacers. They are used on all logic and memory boards except where flow blockers are used. Since the memory die have lesser cooling requirements than the logic die, the memory boards use a special plastic spacer that fits right on top of the die, between the die and the adjacent circuit board. These spacers maintain the correct board-to-board spacing, add mechanical stability and also help to block some of the flow of Fluorinert through the memory modules thereby diverting more of the flow through the logic modules.

Interconnection from one module to another is achieved by means of twisted pair logic connectors which are attached to the module assembly in a semi-rigid fashion. Each circuit board on an outside edge of the module assembly may have twisted pair or flex connectors.

Although the module assembly is the basic removable unit in the electronic section of the CRAY-3, the basic unit of assembly and repair for the module itself is an individual board stack. One of these board stacks can be seen in relationship to the rest of a module in the illustration shown in Figure 95.

FIGURE 95. The Board Stack on a Module.



The module's basic mechanical integrity is provided by the 104 mm circuit boards or plates at the center of the module sandwich. The smaller boards, and hence the stacks, are attached to the plates with power and alignment pins, as well as the complement of twist pin logic jumpers used on each board stack.

Before presenting more detailed discussions of each of the printed circuit boards used in a CRAY-3 module, we will take a few paragraphs to give an overview of how these printed circuit boards are manufactured. The following discussion will be better understood as you read the subsequent discussions and tables describing the different layers of each board.

6.2 Printed Circuit Board Manufacturing

The printed circuit boards for the electronic modules of the CRAY-3 are manufactured by Cray Computer Corporation's own Printed Circuit Board Facility. The circuit boards typically have eight electrically conductive layers separated by seven electrically insulating layers. The two general sizes which are manufactured are 25 mm by 25 mm and 104 mm by 107 mm. Thicknesses range from 0.38 mm to 0.58 mm.

Manufacturing of the circuit boards begins with the lamination of copper foil and dielectric materials. The lay-ups of the appropriate layers are placed between lapped aluminum plates and stacked in openings of a multiple platen vacuum, hydraulic press. The stacks of layers are heated to 400° F using a pre-determined heat rise and pressure profile. This process causes the epoxy/glass insulation material to liquefy and bond to the copper foils on either side. This creates a single layer of material made up of a sandwich of two layers of copper foil and one layer of epoxy/glass insulating material. After three and one-half hours, the layers are removed from the press, trimmed to size and moved to the next processing area for chemical cleaning and photo processing. The thickness tolerances for these sandwiched layers is typically ± 0.00762 mm.

The printed circuit boards use a photo imaging process to define the conductor and pad configurations on the copper foil. This information is supplied by the Artwork Department where these trace lines and pads are drawn using CrayDraw, an in-house CAD program. This digitized data is used to make the glass masks which are used in the photo imaging process at our Printed Circuit Board Facility. These steps at the PCB facility are performed in a class 100 clean room.

The copper layers are coated with an ultraviolet light sensitive polymer film which is 0.0254 mm thick. A photo mask produced from the glass master is registered to the layer on top of the polymer film. Each photomask contains a negative image of the desired conductor configuration. The registered mask and layer are exposed to intense ultraviolet light which causes a linking reaction of the portion of polymer which is not masked. This process is repeated for the appropriate number of layers. After exposure, the layers are sent to the wet processing area.

Here the layers are processed through a spray developer which removes unexposed photoresist. Next, the layers are conveyed through a cupric chloride etcher which removes copper foil that is not covered by the configured photoresist. Finally, the remaining photoresist is stripped off leaving copper conductors matching the original patterns on the photomask.

At this point, layers are microscopically inspected for defects before lamination into boards. Limited repair or rework is also performed by the inspectors. Resistor layers are tested for correct ohmic values using precise automated electrical testers.

Tested and inspected layers return to the lamination area to be assembled into panels utilizing a similar vacuum/hydraulic press process. These panels contain six internal layers of conductors. To provide vertical interconnection of the horizontal inner layers, holes will be drilled through termination pads at each end of the horizontal conductors.

Panels received in the drilling area are registered and punched to within 0.0127 mm on an optical punch system. These punched holes are used to align the panel on the drill tooling plates. Panel alignment of ± 0.00508 mm prior to drilling is accomplished by drilling test coupons and adjusting four independent tooling plates on the drill table.

An average of 24,000 holes are drilled in each panel. Hole sizes range from 0.150 to 0.406 mm, with the majority of holes being drilled at 0.170 mm in diameter. Hole position tolerances are ± 0.00762 mm or less. Following drilling, the drilled panels are ultrasonically cleaned, inspected and moved to the electroless copper plating area.

After chemical pre-treatment, the electroless copper process deposits a 0.001 mm thick "seed" layer of copper on the walls of the drilled holes. The seed layer provides a conductor for subsequent electroplating operations which then provide the conductive path between the inner horizontal conductors and the vertical plated conductors.

The electroless line is a series of chemical tanks housed in an environmentally controlled structure approximately 50 feet long and nine feet high. Panels are processed and transported via a computer controlled hoist system. Upon completion of the electroless plating, the panels are inspected and moved to the clean room for photo imaging of the external plating pattern. Panels are then spray developed and inspected prior to final plating.

In the final plating process, the holes and defined surface patterns are plated with 0.005 mm of copper, 0.005 mm of nickel and 0.001 mm of gold. This plating provides vertical conductors for the board vias, and bonding surface metal for the integrated circuits' gold leads and the twist pin jumpers.

After plating, the photoresist pattern is stripped. The copper remaining around the plated pattern is etched away. Panels are then moved to the saw area for mechanical processing and inspection.

The individual printed circuit boards are separated from the panel using a diamond grit saw blade on an automated dicing saw (the same type of machine

used to dice the GaAs wafers). The sawn boards are inspected and measured for size and thickness tolerances of ± 0.025 mm.

An electrical test follows the mechanical inspection. The DITMCO “bed of nails” tester has a grid of 0.229 mm diameter spring probes which contact every possible plated hole location. The test head for the 104 mm boards has approximately 40,000 probes. The 25 mm head has 5,200 test points in a staggered over/under configuration. All circuit paths and possible open or shorted conditions are tested on each circuit board. Finally, electrically good boards are inspected using microscopes prior to delivery to Module Assembly.

6.3 The 25 mm Circuit Boards

A module assembly contains 64 module circuit boards. These boards are arranged in 16 stacks of four boards each as shown in Figures 92 and 94. Two of the four boards in each stack are above the logic plates and two boards are below the logic plates. Board sizes are 25.27 mm x 25.27 mm for inside locations, 25.27 mm x 28.07 mm for side locations and 28.07 mm x 28.07 mm for corner locations on a module. Board stacks are positioned on 25.4 mm centers.

The module circuit boards are assembled as multi-layer printed circuit boards. There are six inner layers of metal circuits and two exterior layers. Four of the inner layers are for power distribution. These layers are solid metal except for relief holes. The other two inner layers are for logic circuit trace lines. The exterior layers have pads surrounding each plated hole in the circuit board. Exterior surfaces of metal are treated with 0.005 mm copper, 0.005 mm nickel and 0.001 mm gold plating. There are three sizes of plated holes in a module circuit board as shown in Table 10.

TABLE 10. Plated Holes In a Module Circuit Board.

Hole Type	Drilled	Plated	Pads	Relief Holes
Power Holes	16 Mil Bit	15.1 Mil Finished Hole	20 Mils Across Flats	20 Mils Across Flats
Jumper Holes	6.7 Mil Bit	5.8 Mil Finished Hole	9 Mils Across Flats	10 Mils Across Flats
Via Holes	5.9 Mil Bit	5.0 Mil Finished Hole	9 Mils Across Flats	10 Mils Across Flats
IC Pin Holes	5.9 Mil Bit	5.0 Mil Finished Hole	9 Mils Across Flats	10 Mils Across Flats
Hole Position Accuracy is ± 0.3 Mils (thousandths of an inch).				

FIGURE 96. Plated Hole In a Logic Circuit Board.




Table 11 describes the layers in a circuit board by showing what each layer is used for as well listing the specifications of each layer.

TABLE 11. Layers In a CRAY-3 Logic Circuit Board.

Layers	Usage	Specifications
Layer 1	Electrical Pads	Combined plating thickness of 0.6 mils.
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 2	Horizontal Logic Circuit Traces	0.36 mils (1/4 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 3	VG Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 4	VP Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 5	VG Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 6	VN Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 7	Vertical Logic Circuit Traces	0.36 mils (1/4 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 8	Electrical Pads	Combined plating thickness of 0.6 mils.
Total Board Thickness = 15.8 mils (thousandths of an inch).		

A plan view of a printed circuit board is shown in Figure 97. This particular example shows a corner logic board. Remember, the corner circuit boards are the largest of the three sizes used. The three larger size jumper holes near the corners of the board are for the power and alignment pins. The other small jumper holes shown throughout the board are for logic and mechanical jumpers. Mechanical jumpers are logic twist pins that are not carrying any signal other than ground. They help improve ground plane conditions as well as add extra mechanical stability to a board stack.

FIGURE 97. P Stack Logic Circuit Board.



The holes for the gold leads from the die can be seen directly around the square areas empty of any jumper holes. Since the board shown in Figure 97 mounts at the corner of a module this P stack logic board has wirebond jumper holes

along the bottom and right sides of the board. The wirebond jumper holes are the six rows of holes at the edges of the board. The small twisted-pair wires which attach to the 2 by 22 connectors are soldered to the two wirebond jumper holes closest to the edge. The small area, empty of jumper holes, in the center of the board is reserved for labeling of the board. The signal trace lines cannot be seen in this view since they are on inside layers.

FIGURE 98. Corner of a Circuit Board with Trace Lines.




Figure 98 gives a view of one of the inside trace layers of a circuit board. These trace lines are drawn in the Artwork Department using CrayDraw. The circuit

board signal traces are drawn on two layers, one having predominantly vertical traces and one having predominantly horizontal traces. Where a connection is needed between the two layers, a via is placed. Vias are placed off the normal logic jumper grid. Seven vias can be seen in the die area of Figure 98 (identified by the letter "V" inside the pad). A via is a plated through hole which only connects the two trace layers. No twist pin jumper passes through a via. The electrical connection between the layers is made only by the plated barrel of the via hole.

The Artwork Department draws all of the trace lines on every circuit board using a Board Stack Composite which lists all the point-to-point connections in a simple text code. This code can be read by computer programs which automatically route circuit boards. However, even those boards which can be fully routed automatically require hand touch-up by Artwork Technicians.

The Artwork Technicians also employ several computer test programs which will scan the digitized data and alert them to any open circuits, short circuits, spacing that doesn't meet design rules and trace line lengths that may be too long. When the drawing is completed on an entire stack of boards, final test programs are run to be certain that all the connections that are drawn match the intentions of the designer.

In addition to the logic circuit boards, there are two other 25 mm circuit board types used in the CRAY-3. These are the 256K memory boards and the One Meg memory boards. These boards have the same dimensions as the logic boards but differ slightly in the construction of their layers as well as in their jumper layout. The different layer arrangement is necessitated by the fact that the silicon memory die require different voltages than do the GaAs die. The different jumper hole arrangement is required because the silicon memory die have different pin-outs and sizes than the GaAs die. Table 12 shows the layers in a memory circuit board. Figure 99 gives a plan view of a 256K memory circuit board which holds nine die. Figure 100 shows a plan view of a One Meg memory circuit boards which holds six die.

TABLE 12. Layers In a CRAY-3 Memory Circuit Board.

Layers	Usage	Specifications
Layer 1	Electrical Pads	Combined plating thickness of 0.6 mils.
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 2	Horizontal Logic Circuit Traces	0.36 mils (1/4 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 3	VM Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 4	VP Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 5	VG Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 6	VN Power Layer	0.7 mils (1/2 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 7	Vertical Logic Circuit Traces	0.36 mils (1/4 oz.)
Glass Epoxy	Structural Fill and Dielectric	1.58 mils
Layer 8	Electrical Pads	Combined plating thickness of 0.6 mils.
Total Board Thickness = 15.8 mils (thousandths of an inch).		

FIGURE 99. 256K Memory Circuit Board.





FIGURE 100. One Meg Memory Circuit Board.



6.4 Logic Plates

A module assembly contains two module logic plates. These plates form part of the center section of the module, as viewed from the side. The two logic plates in each module are identical mechanically. They differ only in the jumper hole patterns and the circuit traces in the inner layers. These traces provide the logic signal interconnections between the circuit board stacks. The plan view of a module logic plate is illustrated in Figure 101.

The module logic plates are assembled as multi-layer printed circuit boards. They have six inner layers with metal traces for logic signal paths and two exterior layers with pads surrounding each plated hole. Metal used in the fabrication is primarily copper. As with the smaller printed circuit boards, exterior surfaces of metal are treated with 0.005 mm copper, 0.005 mm nickel and then 0.001 mm gold. The layer characteristics are listed in Table 13.

Power holes are drilled with a 0.406 mm bit and are plated to a 0.384 mm finished hole. Jumper holes are drilled with a 0.170 mm bit and are plated to a 0.147 mm hole. Exterior pads are octagonal in shape with 0.508 mm across flats for the power holes and 0.229 mm across flats for the jumper holes. Jumper holes are located on a 0.508 mm grid. Exterior surfaces are filled with unconnected metal trace lines. These fill line traces are 0.127 mm wide and are centered between the jumper hole pads.

Plating in the jumper holes is selectively connected to a trace on an inner layer of the logic plate. All connections are made in pairs with the jumper pair horizontally oriented on the plate. The jumper to the left is always logic true. The jumper to the right is always logic false. The electrical signal arriving at a jumper pair is transmitted along a pair of co-planar lines to a destination jumper pair. The logic signals on the paired traces are complementary with one carrying the logic true signal and the other carrying the logic false signal. The parameters of the co-planar traces make the path appear as a transmission line of 50 ohm impedance for each signal of the pair. The electrical resistance of a co-planar line is about 0.24 ohms per centimeter. Changes in direction through the logic plate are accomplished by vias through the plate to another layer with co-planar lines in the alternate directions.


The co-planar traces consist of metal traces 0.0711 mm wide and spaced 0.0559 mm apart. Jumper hole pads are 0.203 mm across the flats. The pads are on 0.508 mm centers with a single pair of co-planar lines between the pads. This leaves 0.057 mm between the outer edge of a co-planar line and a jumper pad where that trace passes the pad.

TABLE 13. Mechanical and Electrical Layers in a CRAY-3 Logic Plate.

Layer	Usage	Specifications
Layer 1	Pads and Fill Lines	Combined plating thickness of 0.6 mils. Fill lines are vertical.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 2	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented horizontally.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 3	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented vertically.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 4	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented horizontally.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 5	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented vertically.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 6	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented horizontally.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 7	Electrical Traces	0.5 mils (3/8 oz.) Coplanar lines are oriented vertically.
Glass Epoxy	Structural Fill and Dielectric	1.97 mils
Layer 8	Pads and Fill Lines	Combined plating thickness of 0.6 mils. Fill lines are horizontal.
Total Plate Thickness = 18 mils (thousandths of an inch).		

Figure 101 shows a plan view of a module logic plate. Notice that the jumper holes are the same as the jumper holes for the 16 board stacks that would be associated with the two logic plates of that same module. The holes which can be seen within the areas where the die are placed are vias. One can also see the wirebond jumper holes around the top, right side, and bottom edges of the plate.


FIGURE 101. Module Logic Plate.



As with the smaller printed circuit boards, the signal trace lines cannot be seen since they are on inner layers.

The signal trace lines for the logic plates are drawn in the Artwork Department using CrayDraw. The plates are first auto-routed using a computer program which is run on the CRAY-2. Lines that the auto-router were unable to draw are then added by hand using CrayDraw. Changes which may be needed, as well as padding, are also done by hand. An illustration of the logic plate trace lines is shown in Figure 102.

FIGURE 102. Portion of a Logic Plate with Trace Lines.



Note: Because the resolution is not high enough the trace lines shown above appear to be only one solid line. However, in reality they are paired lines which only separate near the destination jumpers. Only one of each trace line contacts each of the paired jumpers.

Each individual trace line is 0.071 mm with a 0.0559 mm gap between the two lines.

6.5 Resistor Plates

Until the end of 1992 a module assembly contained one resistor plate. This one plate was at the center of the module assembly. This plate provides terminating resistors for impedance matching of output signals. In early 1993 experiments were begun with moving the resistor plate from its center location in the module assembly sandwich to a location on the outside of one of the logic plates to allow for additional cooling effect from the Fluorinert. Further experiments are under way to make the resistor plate into individual resistor boards the same size as the 25 mm printed circuit boards. These individual resistor boards would also be mounted on the outside of the logic plates rather than at the center of the module sandwich to allow for better heat dissipation to the Fluorinert flow. As the outcome of these experiments becomes settled, appropriate revisions will be made to the Hardware Description Manual, particularly as these changes relate to this section.

Module resistor plates are assembled as printed circuit boards. Exterior layers have pads surrounding plated jumper holes and rectangular strips surrounding common grounds of clustered resistors. Three of the seven internal layers each contain 6,912 resistor elements for logic modules, 2,592 resistor elements for 256K memory modules and 2,448 resistor elements for one meg memory modules. The total available resistor elements for logic modules is 20,736; for 256K memory modules it is 7,776 and for one meg memory modules it is 7,344. The other four internal layers are ground planes. The resistor elements are 25 ohms per square nickel phosphorus or nickel chrome composition. Resistor element dimensions are 0.305 mm wide x 0.732 to 0.488 mm long (depending on the material) for the 55 ohm logic module resistors and 1.02 mm wide x 2.44 to 1.63 mm long for the 55 ohm memory module resistors.

TABLE 14. Mechanical and Electrical Layers in a CRAY-3 Resistor Plate.

Layer	Usage	Specifications
Layer 1	Pads and Strips.	Combined plating thickness of 0.6 mils.
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 2	Resistor Layer (Ohmega)	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 3	VG Ground Layer	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 4	VG Ground Layer	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils

Layer	Usage	Specifications
Layer 5	Resistor Layer (Ohmega)	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 6	VG Ground Layer	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 7	VG Ground Layer	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 8	Resistor Layer (Ohmega)	0.38 mils
Glass Epoxy	Structural Fill and Dielectric	2.1 mils
Layer 9	Pads and Strips	Combined plating thickness of 0.6 mils.
Total Plate Thickness =20.66 mils (thousandths of an inch).		

Power feed-through holes are drilled with a 0.41 mm bit and are plated to a 0.383 mm finished hole. Jumper holes are drilled with a 0.170 mm bit and are plated to a 0.147 mm hole. Power blade areas are non-functional but serve as a feed-through for power blade jumpers. Jumper locations are selectively connected to resistor or ground locations on the internal layers.

A plan view of a resistor plate is shown in Figure 103. The white dots in groups of four along the left edge of the plate are holes that allow the power blade jumpers to feed through. Compare the enlarged illustration in Figure 104 with the plan view of the resistor plate in Figure 103 to help you discern the layout of the parts.

FIGURE 103. 256 K Memory Resistor Plate.





FIGURE 104. Resistors on a Logic Module Resistor Plate.



6.6 Power Plates

A module assembly contains two module power plates. These plates form part of the center section of the module. The power plates provide the power distribution from the module power blades to the circuit board stacks.

Module power plates are assembled as multi-layer printed circuit boards. Inner layers of metal are solid except for relief holes. Exterior layers have a pad surrounding each hole through the plate except where electrically isolated 0.170 mm holes allow passage of jumper pins from upper stack locations to lower stack locations. All pads are octagonal shapes. Metal used in the fabrication is primarily copper. Exterior surfaces of metal are treated with copper, nickel and then gold (0.00508 mm Cu, 0.00508 mm Ni, 0.00127 mm Au).

The plan view of a module power plate is illustrated in Figure 105. The small circles represent plated holes for power pin jumpers into the circuit board stacks. The rectangular fields on the left edge of the plate represent arrays of 64 holes each for power pin jumpers into the module power blades. All holes are drilled with a 0.406 mm bit and are plated to a 0.384 mm finished hole. Plated pads are 0.508 mm across flats.


The exterior surfaces of the module power plate are filled with a metal pattern of repeating squares, except where jumper feed through holes exist, in the areas of the plate corresponding to the board stack positions. Each square is 0.305 mm on a side and the squares are spaced 0.203 mm apart. This pattern assists in plating uniformity.

TABLE 15. Mechanical and Electrical Layers in a CRAY-3 Power Plate.

Layer	Usage	Specifications
Layer 1	Pads and Fill Pattern	Combined plating thickness of 0.6 mils.
Glass Epoxy	Structural Fill and Dielectric	1.9 mils
Layer 2	VP Power	2.8 mils (2 oz.) Relief holes 20 mils across flats.
Glass Epoxy	Structural Fill and Dielectric	1.9 mils
Layer 3	VG Power	2.8 mils (2 oz.) Relief holes 20 mils across flats.
Glass Epoxy	Structural Fill and Dielectric	1.9 mils
Layer 4	VN Power	2.8 mils (2 oz.) Relief holes 20 mils across flats.

Layer	Usage	Specifications
Glass Epoxy	Structural Fill and Dielectric	1.9 mils
Layer 5	VM Power	2.8 mils (2 oz.) Relief holes 20 mils across flats.
Glass Epoxy	Structural Fill and Dielectric	1.9 mils
Layer 6	Pads and Fill Pattern	Combined plating thickness of 0.6 mils.
Total Plate Thickness = 22 mils (thousandths of an inch).		

FIGURE 105. 256K Memory Module Power Plate.




6.7 Memory Board Bypass Capacitors.

In early 1993 design changes were planned in the memory boards to allow for the placement of very small bypass capacitors directly on the 25 mm printed circuit boards used for Common Memory. These design changes were being planned to eliminate the high frequency noise that is presently seen on the Vm bus bar during memory activity. This noise is caused when large numbers of memory circuits are being switched on and off at the same time. Although the linear power supplies are rigid enough to handle this switching activity, the inductance of the bus bars delivering this rapidly switching power causes droops in the Vm voltage at the load. By placing the capacitors directly on the memory printed circuit boards, in intimate contact with the board power layers, ample surge current should be provided to the memory die at CRAY-3 operating speeds.

The following data was accurate at the time this revision of the Hardware Description Manual was printed, but the project was still in the design stage and specifications were not all finalized. Twelve surface mount capacitors, mounted in two rows of six alongside opposite top and bottom edges of the SA, SE and SM boards is planned. These capacitors will be 0805 size packages 0.08 inch by 0.05 inch, constructed with Y5V ceramic dielectric. They will deviate from the normal 0805 specification in that they will be a maximum of 0.018 inch thick instead of the usual 0.05 inch. Each capacitor will be 0.1 μ f, with a 16 volt breakdown voltage rating. The capacitors will be mounted to the boards prior to die attach using conductive adhesive or solder.

Figure 106 gives a preliminary view of where the capacitors would be mounted on a 25 mm SE memory board. The die shown in this illustration are four meg Toshiba memory die, which are part of a planned memory enhancement for the CRAY-3 (see Chapter 1, page 18).

FIGURE 106. SM Die and Bypass Capacitors.




6.8 Power Blades

Each module assembly has four module power blades. Each blade is associated with one of the three power supply voltages or ground reference. The blades connect the module power plates to the horizontal power buses, which are a part of the octant assembly structure. These four connections provide the mechanical support for the module as well as the electrical power connections (see Figures 141-144 in Chapter 8).

The module power blades are machined from a high purity copper bar and the resulting part is plated first with nickel and then with gold. The module power plates and termination resistor plate are connected to each module power blade by 64 power pin jumpers. The plates are inserted into a slot on the blades and then secured with the 64 power pins by driving the pins into aligned through-holes in the blades and plates. Bulges in the pins interface with the through-holes in the power plates.


FIGURE 107. Module Power Blade.



The module power blade is tapered and is forced into a rectangular slot in the horizontal bus bar by a screw. Each insertion of the module results in a sliding

deformation of the contact area between the blade and the distribution block. This provides a good electrical contact for the high current flow required for the module. One reason for the large number of jumpers in the connection between the module power plates and the power blades is the potentially large shear forces which may be generated by driving the bus screws. The four screws involved in seating a module are driven in sequence with a torque screwdriver to minimize the amount of shear force between the power blades.

FIGURE 108. Vp Power Distribution In a Module.



The distribution of power and ground throughout a module begins at the power blades. The power blades receive power from the power supplies through their attachment to the horizontal power bus. The blades transfer power to the power plates where the electrical power is distributed to all of the power pins located

at the corners of each printed circuit board which makes up the logic and memory stacks. The power pins in turn distribute the power and ground to the power and ground layers of each printed circuit board where the appropriate gold pins from each integrated circuit package contact the appropriate power and ground layers in the board to transfer the correct voltages and ground reference into the packages themselves. This power distribution is illustrated in Figures 108 through 111.

FIGURE 109. Power Distribution In Logic Modules.









FIGURE 110. Power Distribution in 256K Memory Modules.




Note: On memory modules all the board stacks use memory boards except for D, H, L, and P. These are the control stacks and they use logic boards.

FIGURE 111. Power Distribution in Four Meg Memory Modules.



Note: On memory modules all the board stacks use memory boards except for D, H, L, and P. These are the control stacks.




6.9 Logic Connectors

Each module assembly has a number of logic connectors for electrical communication with other modules in the system. Primary communication is between the left edge of the module and the central wire mat. Secondary communication is local to small groups of modules and is between the top and bottom edges of the modules in the group (see Figures 150 and 151 in Chapter 8). The local communication paths are in the direct fluid flow through the modules.

Electrical communication between modules is over twisted pair wire paths. Recently, experiments have begun using special coaxial cables rather than the individual twisted pair wires. The signals on the pairs of wires are complementary and the characteristic impedance of the wire pair is approximately 60 ohms. The signals leave the module from the edge of a module circuit board. The twisted pair of wires is soldered to plated hole pads on the circuit board. The signal travels over the pair to a semi-rigid logic connector. The wire path between edge of board and logic connector body is approximately 3.81 mm.

FIGURE 112. Logic Connectors Connected to a Corner Logic Board.



Note: This illustration shows Flex cable instead of individual twisted pair wires.


The logic connector configuration is illustrated in Figure 112. There are 44 wires from each board which form 22 twisted pairs. Each wire consists of seven strands of number 44 gauge copper wire with silver plating twisted at 6 twists per inch. Diameter of the wire bundle is 0.152 mm and is equivalent to a 36 gauge solid wire. Insulation is Teflon with a 0.051 mm wall thickness.

The connector body is a molded plastic part with strip line contacts. The plug contacts consist of a formed beryllium-copper alloy spring of two leaves which is then hard gold plated. The receptacle contacts consist of beryllium-copper barrels which are gold plated.

One logic connector is associated with the upper (or lower) half of a board stack in a module. The length of the logic connector is slightly less than half the board height. The logic connector thickness is slightly less than the module thickness. The left edge of the module accommodates eight logic connectors spaced at a 3.8 mm distance in a semi-rigid structure. Those modules which have secondary logic connectors for local communication use the top and bottom edges of the module in a similar manner.

Recently, a new type of wire has been used between the stripline connector and the wirebonds on the edge of the printed circuit boards (see Figure 113). Instead of using individual twisted-pair wires, a piece of flex circuit containing 44 traces has been used. The use of this flex circuit has allowed for a simplified soldering process as well as the elimination of several parts, including the strain reliefs and corner spacers. An epoxy glue is now used to permanently fasten the flex circuit to the boards after the traces have been soldered in place. This method of attaching the logic connectors to the boards holds the promise of improved mechanical, and hence electrical, resiliency. It has also reduced the module assembly time by approximately two hours per module.

FIGURE 113. Module Logic Connector with Flex Circuit.



Now that we have a basic understanding of module structure we can go on to see how a module is assembled, tested and repaired in Chapter 7. These are important operations in building working machines and demand a large percentage of time from many people in manufacturing and module test.

Module Assembly, Testing and Repair

Much of the day-to-day activity in building working versions of the CRAY-3 revolves around the modules. New modules are built and released for testing after final assembly. Modules in the test cycle are tested, problems diagnosed, and then sent back for disassembly and repair if needed. During this entire cycle of build, test, and repair the failures are analyzed so that any inherent design or manufacturing weaknesses may be detected and corrected. In this chapter we review this process beginning with the building of a new module.


7.1 Final Assembly of the Module

Module Assembly is the process of building a module consisting of sixty-four circuit board assemblies grouped into sixteen four-board stacks. The stacks are placed onto a power plate assembly consisting of the resistor plate, the two power plates, and the two logic plates. After stacking, each stack is inserted with power and twist pins that hold the module together mechanically as well as make the necessary z-axis intra-module electrical connections for power and logic signals.

The process of building a new module begins by ascertaining the needs of Module Test for particular modules. This information is made known to Production Dispatch where build schedules are implemented. Inventories of parts are checked to make sure there are enough of the correct parts on hand to

begin the build process for a module. The basic steps involved in building a module assembly is shown in Figure 114.

FIGURE 114. Module Build Flow Chart.




The assembly of a module takes from three to eight days, depending on the complexity of the module. The process begins by choosing the appropriate die types needed for each one of the 64 boards of that particular module and then placing those die in kits that are used by Die Attach. This first step in the assembly of a specific module is called “kitting”.

7.1.1 Kitting

A specific type of module actually begins to take shape in inventory with the kitting of the unique type of die needed for each of the 64 boards that makes up that module. The kitting process takes specific die types from inventory and places them together in waffle packs which represent the die make-up of each board on the module. At this point those die are removed from inventory listings and made part of a named module.

FIGURE 115. The Kitting of Die for a Module.




After all the die have been kitted for a module, the kitted waffle packs are delivered to Die Attach where the correct bonded die are attached to the printed circuit boards which have come from Board Solder.

7.1.2 Board Soldering

Before the printed circuit boards are delivered to Die Attach they are cleaned and boxed for a particular module assembly. The boards which will be used on the edge positions of the module (A, B, C, D, H, L, M, N, O and P Stacks) will have very small, twisted pair wires soldered to the wirebond areas. These tasks are performed in Board Solder.

FIGURE 116. Small Wires Soldered to Circuit Boards.



Twenty-two, twisted pair wires attached to 2 by 22 molded connectors arrive from a vendor. The free ends of the connector wires are prepared in Board Solder by being cut to the specific lengths which have been listed in a Wire Mat document. They are then untwisted to have their individual lead ends stripped and tinned. A laser strip operation is used to cut the teflon insulation back from the conductors. The exposed leads are then tinned, inspected, and tinned again. The connector is ultrasonically cleaned to remove all traces of flux residue. The prepared wires are then inserted into a fixture along with the corresponding circuit board. After inspecting for proper alignment with the wirebond pads on the boards, the prepared wires are soldered onto the board using solder quick tape at a Pulsed Reflow station. All boards, except for the logic stacks on memory modules, will have small wire strain reliefs attached at this point using an epoxy glue. The boards are then cleaned and all the soldered areas inspected. After final inspection the boards are delivered to Die Attach.

A new part is currently being tested which replaces the individual twisted-pair wires with flex cable. This piece of flex cable comes with 44 electrical traces imbedded between layers of flexible Mylar. The flex cable is soldered to the wirebonds on the printed circuit boards using a special Re-Flow Fixture illustrated in Figure 117. Strain reliefs are not used with the flex cable connectors since an epoxy glue is used to permanently connect the Mylar material to the circuit boards after the small traces are soldered to the wirebonds. Two flex connectors attached to a printed circuit board are shown in Figure 118.

The flex connectors have shown great promise. They allow for a simplified soldering process, eliminate the need for strain reliefs and corner spacers, should give better mechanical endurance and have reduced module assembly time by about two hours for each module.

FIGURE 117. Flex Connector Re-Flow Fixture.





FIGURE 118. Flex Circuits Attached to Board Wirebonds.



7.1.3 Die Attach

As we have seen, the assembly of a module begins with the kitting of the appropriate die population and the preparation of the printed circuit boards. Printed circuit boards which are to be on the edge of a module will have been prepared with twisted pair wires soldered to the wirebonds on the edges of those boards at Board Solder. These two component parts, the kitted die and the soldered and cleaned boards, come together at Die Attach.


At Die Attach 16 logic package die with bonded leads are attached to the individual printed circuit boards. Memory circuit boards have varying die populations and are treated separately. However, the process of attachment for memory circuit boards and logic circuit boards is quite similar. The differences

between memory board assemblies and logic board assemblies are in the differing die population, number of leads, die sizes and die positioning.

Correct board and die orientation, straightness of the bonded die leads, and machining tolerances of the die attach fixture are critical for proper assembly. Circuit boards one and two receive the die from the layer eight side of the board. Circuit boards three and four receive the die from the layer one side of the board. This is because the die are always on the inside of the board facing towards the center of the module assembly even though boards one and two are on the opposite side of the plate assembly from boards three and four.

The die are moved from the waffle pack, in proper formation and proper individual orientation, to the die attach fixture using a vacuum pen. The operator is guided by documentation that travels with the kitted die. The documentation is produced directly from the Board Stack Composite which was used by the Artwork Technicians to draw the board traces. The fixture is a finely machined stainless steel waffle pack block in which the die lie with the leads facing upwards in appropriate orientation to receive the printed circuit board's drilled and plated die lead holes. The die are held in individual pockets within the die attach fixture by means of vacuum. The die and the fixture are checked for proper orientation and for lead straightness under a microscope. Stainless steel guide pins are placed on diagonals across the fixture. Guided by the guide pins on the fixture, the printed circuit board is lowered over the leads of the individual die. Then the board is physically clamped in place over the die. At this point in the process an individual lead from a die is flush with the opposite surface of the board (see Figure 119).

FIGURE 119. Die Attach Stamping Fixture.



The complete fixture holding the board and the die is now placed in a stamper. The stamper is a modified Hughes bonding machine which forces a shaped hardened tool to compress the gold lead into the barrel of the printed circuit board. The compression is controlled by using a pre-determined force. The lead is compressed to a depth of about 0.178 mm into the barrel forming a plug. Notice the illustration of a stamped lead in Figure 119.


After stamping, the board assembly is inspected for the presence of any cracked die, any imperfections on the backside surface of the die, the quality of the lead stamping shown on the opposite side of the board and the integrity of the strain reliefs on boards with wirebonds and twisted pair wires.

When inspection is complete, the finished circuit board assembly is placed in an anti-static container and delivered to Module Assembly for stacking. There it will become part of a complete module assembly.

7.1.4 Twist Pin Manufacturing

Essential to the mechanical and electrical performance of a CRAY-3 module are the thousands of twist pins used in the assembly of each module. These twist pins are made from gold-plated beryllium-copper wire. The twist pins used to transfer logic signals from one board to another are made from a seven strand, 0.122 mm wire. They are often referred to as “jumpers”. The twist pins used to transfer electrical power and ground from the power plates to the circuit boards are made from a 19 strand, 0.318 mm wire. Both types of twist pins are shown in Figure 120.

FIGURE 120. Twist Pin Illustrations.



Both the twist pins used for logic jumpers and the twist pins used for power pins are made using Hughes Twist Pin machines. The beryllium-copper wire is

placed in the machines on spools and threaded into a twister arm which untwists the wire at pre-determined distances to produce a bulge. The machine is equipped with a Class IV Neodymium Glass laser which cuts the pins to their specified lengths. The laser also makes small welds at both ends of each twist pin. The machine operation is controlled by a computer which reads the instruction software from floppy diskettes. Once the wire spools have been loaded and the wire properly threaded, production of the entered quantity of pins is automatic.

After the twist pins have been properly formed, cut and welded by the Twist Pin machines, they are placed into Pyrex containers according to machine and lot number. They are then heat-treated in an inert gas oven to enhance the bulge strength. The heat treatment process is completely automatic. During the first thirty minutes a nitrogen purge of the chamber is done at room temperature. Following the purge cycle, the oven temperature is gradually raised to 550° F and maintained for two hours and 45 minutes. The pins are then removed for inspection.

The twist and power pins must be inspected regularly to insure that they are being made within specifications. Pins that do not meet the proper specifications could cause unreliable logic and power operation in completed modules. Bad pins could also cause electrical shorts and barrel damage in printed circuit boards. The pins are inspected randomly. They are visually checked for proper cuts at both the tail end (nearest the last bulge of the pin) and the leader end (that portion of the pin which is first inserted through the layers of the module). Other non-dimensional defects which cause rejection are mushrooming of a weld, fraying (incomplete welding), club ends (off-centered welds) and hooks (hook-like formations at the leader end of the pin). Some of these defects are illustrated in Figure 121.

FIGURE 121. Twist Pin Defects.

Ideal Tail Weld



Ideal Leader Cone



Tail Defects



Mushroom



Frayed



Club




Hook




Pointed

Leader Defects



Hook



Club



Unacceptable Cone



Worst Case Acceptable Cone

A sampling of pins is also measured under a microscope for overall length, width of the bulges and proper curvature of the bulges. Finally, the pins are counted, boxed and labeled for delivery to Inventory where they await their use in the pinning of Module Assemblies.

7.1.5 Stacking the Module Assembly

All of the separate parts of a module come together to make up a complete module when a module assembly is stacked. The final electrical and mechanical connections for the module are completed when the stacked module assembly is pinned.


The parts used to stack a module assembly include 64 board assemblies from Die Attach; the parts that make up the plate assembly, namely one resistor plate, two power plates and two logic plates; guide pins; center and corner spacers (in the case of memory boards, a special one-piece spacer is used); adhesive flow blockers; and Ultem overlays for each stack. Parts are inspected for any damage or abnormalities before and during the stacking process.

A module is assembled in two stages. The first stage combines the circuit board stacks one at a time with the central group of five plates in a process of stacking on a specially designed fixture. The second stage connects the module power blades to the longer edge of the stacked module.

The module assembly is stacked on a stacking fixture one layer at a time, beginning with the Ultem overlays on the top of a module (which are next to the board ones of each stack) and proceeding through all the constituent layers of a module, ending with the Ultem overlays on the bottom (next to the board fours of each stack).


There are two types of stacking fixtures used: the Universal Stacking Fixture and the 16-Up Stacking Fixture. The basic difference between the two fixtures is that the Universal Stacking Fixture has toggle switches which allow the fixture guide pins to be raised or lowered by individual stacks. This allows for any combination of stacking, from a single stack to a complete module. The 16-Up Stacking Fixture's guide pins can only be raised and lowered for all 16 stacks at the same time (see Figure 122).

FIGURE 122. Stacking a Module.



Once all layers of the module have been stacked, the fixture guide pins can be retracted leaving the permanent guide pins, which were inserted during the stacking process (four per stack), holding the module assembly together.

FIGURE 123. Illustration of Stacking Layers.



Before leaving the stacking area, the stacked module will have the power blades attached to the module assembly. A label identifying the module with a four character code will be applied to one of the outboard power blades. The power blades are attached to the protruding power plate using from four to six power pins per blade. These are inserted by hand. The power pins go through the rim of each blade and through the edges of the power plates. The power plates and the central terminating resistor plate are 0.381 mm longer in one dimension to allow for the attachment of the power blades (see Figure 105). Later, the remaining power pins (64 per blade, total) will be inserted by one of the pin insertion machines.

After a final inspection of the stacked module, it is placed in a Power Pin Palette. The module assembly, in its Power Pin Palette, now has the power pins for each stack inserted by a Semi-Automatic Pin Insertion machine specially modified for inserting the larger power pins. When the 19-strand power pins have been fully inserted in the stacks, the remaining power pins for the power

blades are inserted. After all the power pins have been inserted into the power blades, the module will be placed in a regular Pinning Palette if destined for the Semi-Automatic Pin Insertion Machines so that the logic jumper pins can all be inserted. If the module is to be pinned using the Automatic Pin Insertion machine a palette is not used. Notice that during stacking the only pins that are inserted in the stacks are four guide pins on the inside corners of each board stack (64 pins altogether). During module pinning all of the power pins and logic pins will be inserted (this amounts to about 14,000 pins for each logic module).


7.1.6 Pinning the Module Assembly

The stacked Module Assemblies are pinned using Hughes Semi-Automatic and Automatic Pin Insertion machines. The Semi-Automatic machine is a modified Hughes ball bonding machine. It consists of a puller and a cutter mechanism which are operated pneumatically. The process has been recently modified to include a brake on the cutter. The Semi-Automatic Pin Insertion machine allows the operator to insert pre-formed logic pins into the board stacks in a pre-arranged sequence, guided by a computer program. The pre-formed pin has a leader of the same length as the thickness of the module which allows the pin to be dropped through from the top and then grasped and pulled through from the bottom so as to seat the bulges at each circuit board jumper hole barrel (see Figures 93, 96 and 120). During the pinning operation, the stack is held securely by the stainless steel guide pins in the four corners of each board. The stainless steel guide pins also hold the outer overlays together by means of an interference fit.

The module is loaded into a module pallet which suspends the module with six metal fingers—two on each side and two along the rear edge. The power blades rest inside a power blade retainer. The module pallet slides into two rails on the Semi-Automatic machine. These rails position the module at the proper vertical height above the puller and cutter mechanism. The rails that accept the module pallet are attached to a motor driven table that can position the module horizontally so that any twist pin hole in the module is directly over the puller.

The operator selects a computer program for any one of the 16 stacks that are to be pinned. Only one stack is selected at a time. This program contains all of the coordinates for each twist-pin hole in the stack that needs to be pinned. The machine then positions each hole to be pinned over the puller by moving the table holding the module with a step motor. It is guided by the programmed coordinates as well as pattern recognition of the holes themselves. Once the machine has properly located a hole to be pinned, it waits for the operator to insert a twist pin.

FIGURE 124. Twist Pin Seating.



Using a tweezer, the operator inserts a twist pin, leader end first, into the hole as deep as the first bulge. The leader protrudes below the module by about 1.52 mm. When the operator has properly inserted the twist pin, a foot switch is depressed and the puller-cutter cycle begins.

Upon activation of the foot switch, the puller-up air solenoid is activated causing the dual-acting puller piston to drive the puller up to a hard stop. This

places the puller blades to within 1.27 mm of the bottom of the module. Next, the puller-clamp air solenoid is energized causing the puller blades to close and clamp onto the twist pin leader. Then the puller-up air solenoid is de-energized and the puller-down air solenoid is energized forcing the dual-acting puller piston back down to its lower hard stop position. A hydraulic damper enables the puller to retract at a slow, constant speed. The lower hard stop position is adjustable so that the proper pin seating can be achieved (see Figure 124).

Following the puller cycle, the cutter-up air solenoid is turned on forcing the cutter-up piston to drive the cutters up to within 1.27 mm of the module. The cutters sit to the right of the puller and must be brought on-line in order to cut the pin leader. So the on-line air solenoid is then activated, bringing the cutters over to the twist pin. The puller, now retracted, sits directly below the cutters. The cutter-lift air solenoid energizes and causes the cutters to lift at a constant rate towards the module. As the cutters approach to within 0.254 mm of the module a sensor is activated. This sensor applies a brake to stop the blades at a point between 0.254 mm and flush with the module. The cutter cut-off air solenoid is now activated causing the cutters to come together on either side of the twist pin leader to cut it off. Finally, the cutters open and the assembly retracts and lowers as the cutter cut-off, lift, on-line and up, air solenoids all de-energize.

Once one pin has been inserted and the leader cut off, the table holding the module positions the module so that the next hole to be pinned is directly over the puller. Then the cycle begins again. The total cycle time for the pinning process of one twist-pin is about 1.5 seconds, plus the time required by the operator to insert the pin and depress the foot switch. Approximately 2,400 pins can be inserted per shift for each of the Semi-Automatic pin insertion machines.

The Automatic Pin Insertion machine uses a leader-less twist pin. Unlike the twist pin used with the Semi-Automatic Pin Insertion machines, the leader-less twist pin is symmetrical. Therefore, a pin may be inserted into a jumper hole in either pin direction. This eliminates the steps required for proper pin orientation (done by the operator in semi-automatic insertion) and the cutting of the pin leader. Specifications for the leader-less twist pin are identical to the pins used for semi-automatic insertion except for the absence of the leader and a slightly smaller bulge size. The smaller bulge size is necessitated by the differences in the insertion dynamics of the two processes.

The automatic pin insertion process begins by loading a specially designed cartridge with twist pins. Twist pins are placed in a vibratory bowl-feeder which automatically separates a mass of twist pins into individual twist pins. These individual twist pins can then be transported in single file by a venturi nozzle to an empty location in the special 24,000 hole cartridge used by the Automatic Pin Insertion machine. The cartridge is mounted on a computer controlled x-y stage. A photo-optical sensor at the bowl feeder is triggered by

each passing twist pin. This signals the x-y stage to move the cartridge to the next empty hole. It also starts a timed air blow which prevents successive pins from coming too quickly. The loading process continues until all 24,000 holes in the cartridge are filled with twist pins. The 24,000 hole cartridge is automatically loaded in less than seven hours at the rate of about 3,600 pins per hour.

The Automatic Pin Insertion machine consists of a custom designed twist pin insertion head on an x-y-z axes robotic assembly machine base. A module is loaded onto a back-lighted fixture. Pattern recognition is used to locate the correct module position. Individual twist pin hole locations are determined from an x-y coordinate map programmed into the control computer. The machine's x-y positioning is so accurate that pattern recognition of the individual holes is not required. This helps increase the machine's speed over that of the Semi-Automatic machines.

The insertion head has a counter-balanced foot assembly which touches down on the module with a force of 15 to 20 grams. During the foot touchdown, the height of the module is measured so that the exact insertion height can be determined. A twist pin is blown from the cartridge into a guide region on the foot. When the machine photo-optically senses a twist pin in the foot guide, a punch is rotated into position and then driven through the foot guide until the programmed insertion height is reached. The machine monitors insertion forces. It will stop an insertion if it senses excessive insertion force. Data can be collected for each insertion or for only those insertions involving error conditions. Insertion forces average about 125 grams. The Automatic Pin Insertion machine can insert about 2,000 pins per hour.

7.2 Module Testing

Completed modules are delivered to Module Test where they are connected to pulse-power test stations which can test modules singly or in groups of up to two processors with memory. Modules tested under pulse power do not require liquid cooling since they are pulsed on for only very short periods of time. Modules tested under full power must be liquid cooled.


7.2.1 Pulse-Power Testing

Pulse-power testing enables the basic electrical functionality and mechanical integrity of CRAY-3 modules to be verified prior to installation in a full-power, liquid-cooled environment. This type of testing is safer for the initial power-on of the modules. The risk of catastrophic failure in a full-power environment is reduced by first testing a module under pulse power where potentially serious problems like short circuits or cracked die can be detected and addressed without causing further damage to the module or adjacent modules on the horizontal power bus. It also allows for easier module installation and removal, since you do not need to disassemble and drain a system tank and then refill it. Further, pulse-power testing allows for direct electrical probing of the circuit jumpers in a module, which cannot be done under full-power testing. This ability provides for far more detailed electrical troubleshooting and analysis than is possible in a system tank environment. It therefore serves as a valuable failure analysis tool for issues identified in full-power system testing.

There are two basic types of pulse-power test station configurations. One type of test station is designed to test a single module or a single processor. The other type of test station can accommodate a Foreground Processor (one E module), two Background Processors (two each of an A, B, C and D module), and 32 memory modules.

Single module and one-processor test stations are configured to allow for the testing of a single module type. The hardware requirements vary depending on the module type to be tested; i.e., whether it is a Foreground Processor, Background Processor, memory module or I/O module. But in all cases the station is set up so that the module to be tested is laid flat on the test bed and screwed into the power bus. This positioning allows access to the circuit interconnect jumpers with oscilloscope probes for failure diagnosis and analysis.

FIGURE 125. Pulse-Power Test Station.



Two-processor test stations are configured to allow testing of a complete two-processor system without installation in a liquid-cooled environment. The modules are installed in a vertical position similar to their positioning in an octant cabinet. System interconnect wire harnesses between the Foreground Processor, Background Processors and the memory modules are the same as those used in a full-power system. Use of contact probes with an oscilloscope is impractical with this arrangement of modules, but inductive probes can be used on the wire harness to observe the performance of signal paths.


TABLE 16. Equipment Used in Pulse-Power Testing.

Equipment	Function In Testing Procedures
Pulse-Power Supply.	Provides the required voltages (V_m , V_p , V_n) and ground (V_g) to the modules under test at a safe duty cycle.
Test Bed.	Provides for the proper mechanical and electrical positioning of the modules under test. The test bed consists of a dielectric block with aluminum and gold-plated copper parts for the power bus assemblies. Cables connect the power supply to the power bus.
Macintosh Computer.	Runs the test software and provides the user interface for operating the power supplies.
Foreground Processor.	An E Module serves as the interface between the modules under test and the Macintosh. It receives test control and data signals from the Macintosh via the buffer board and transmits them to the modules being tested. Signals from the modules under test are then fed back to the Macintosh through the E Module via the buffer board.
Buffer Board.	Serves as a data way station between the Foreground Processor and the Macintosh.
Oscilloscope.	Used for contact and inductive probing of high speed test signals.
Digital Multimeter.	The DMM is used to measure the DC characteristics of circuits under test including resistance, continuity, voltage and current.
Frequency Generator.	Used to provide a clock signal for the modules under test. The nominal frequency used is 500 Mhz.
Hand Tools.	Adjustable torque screw driver. Connect and disconnect tools for the module wireharness connectors. Test probes.

The power supplies used in pulse-power testing are custom designed and manufactured at Cray Computer Corporation. They are linear power supplies with the addition of pulse control, voltage regulation, and protection circuitry. The power supplies operate from 110 volts AC, 60 Hz power. This voltage is applied to step-down transformers and rectified. The rectified voltage is then filtered and stored until the energy is needed. The power supplies have a total storage capacity of 1.62 farads for V_p , 0.54 farads for V_n , and 0.23 farads for V_m . The switching element is a Metal Oxide Semiconductor Field Effect Transistor module. Seven of these modules are connected in parallel to switch the V_p voltage of +3.3 volts at up to 2,100 amperes. One MOSFET module is used to switch the V_n voltage of -1.2 volts at up to 600 amperes and two are used to switch the V_m voltage of +5.0 volts at up to 800 amperes. The MOSFET modules are themselves switched on and off by the pulse control circuits in the power supply. When the voltages are switched on energy is drawn from the large capacitors under control of the voltage regulators. After the power-on pulse is delivered, the regulators are switched off again which in

turn switches off the MOSFET modules. During the time that the power is off, the energy in the capacitors is replenished by the transformers and rectifiers.

FIGURE 126. Pulse-Power Supply Block Diagram.



The voltages are regulated and can be adjusted individually from a remote control box. This allows the CRAY-3 modules under test to be tested at voltages different from their ideal voltage in order to find their operating margins. The voltages can be checked by probing a module power pin with an Oscilloscope or by reading the listed voltages in a window on the test station Macintosh monitor.

While power is applied to the modules under test, the protection circuits in the power supply are monitoring each voltage for over and under voltage conditions as well as for over current conditions. If any of these conditions occurs, the power supply will disable successive power supply output pulses. The detected fault is then indicated on the remote control box with the use of Light Emitting Diodes. To clear the fault indication a reset button may be pressed on the remote control box. If the fault condition still exists, it will be detected again and indicated with an LED on the next power pulse.


The maximum on-time of a power pulse is kept to five milliseconds or less by another protection circuit within the power supply. This prevents a catastrophic failure of both the power supply and the modules under test in the event the Macintosh should apply an enable pulse cycle which was too long. The power supplies are not protected against excessive duty cycle. Therefore, proper care must be exercised by the test person to operate the system within reasonable limits.

When a module is to be tested, a Foreground Processor module is secured in the test bed and connections are made to the clock source and console

communications cables. The module to be tested is secured in the proper test bed power bus. All bus screws are tightened using an adjustable torque screw driver. The appropriate interconnect harness is attached to connect the module to the Foreground Processor and any other test bed modules that may also be present. Applicable test software is selected from the test application menu on the Macintosh, the desired power pulse on-time and duty cycle are selected from the menu bar and the tests are run.

The Macintosh computer first sends a Transistor-to-Transistor Logic level pulse, corresponding to the on-time and duty cycle selected, to the pulse-power supply. When this pulse is received, the power supply applies power to the E module and the modules under test. After a set delay time, which allows the applied power to stabilize, a Dead Start signal is brought to a 1 (low voltage level) and applied to the E module. During this time, the E module local memory is loaded with the test data from the buffer board. Once the test data has been loaded, the Dead Start signal is brought to a 0 (high voltage level) and the E module takes over to run the tests at the clock speed selected at the RF generator. The clock speed is generally adjusted in 1 mHz steps (or greater) to test the clock speed operating margins of the modules under test.

FIGURE 127. Time Chart for Pulse-Power Testing.



When the test results have been received by the buffer board, a signal is sent to the Macintosh to turn off the power supply pulse. The results of the tests are then sent back to the Macintosh to be displayed on the monitor for the test person to examine. The process is started again after a pause time calculated by the duty cycle previously selected. Figure 127 illustrates these timing relationships.

In general, errors are detected when the expected data for a given test packet differs from the data received from the module under test. The application will report the error on the Macintosh monitor. Errors may be logged and stored to a report file using the Failure Consolidation Program. Failures are isolated and verified by using the oscilloscope or digital multimeter. The failures are entered into an open items list located in the front of the test module's individual log book. In-depth descriptions of procedures and failure analysis are recorded in the body of the log book. Upon completion of testing, modules are either sent on to full-power testing or returned to repair, if errors need to be addressed there.

All of the test software used in Module Test is generated in house, primarily by the Engineering Software group. Since the Foreground Processor is the link between the Macintosh and the modules under test, all test code must ultimately be handled using the foreground instruction set.

TABLE 17. Software Used in Pulse-Power Testing.

Program Name	Functions Performed
Foreground Packet Generator.	FPG enables the user to generate test packets based on a Foreground Processor source code called a skeleton. FPG makes the manipulation of data and control patterns quite simple. A large number of test packets may be generated in a short period of time. The test packets generated are called MacPacs.
Background Packet Generator.	BPG enables the user to generate test packets based on a Background Processor source code. These packets are run within a Foreground Processor program called a monitor. The test packets generated are called BacPacs.
Failure Consolidation Package (formerly Failure Analysis Package).	FCP logs and stores test errors. The FCP digests all errors received from the test platform, consolidates and then analyzes these errors. The output from FCP will normally give the user a pattern of fault which in turn may point to a failing piece of hardware.
System Test Interface (formerly System Test Platform).	STI is the user interface used for selecting and running all the test software.

Modules sent to repair undergo a variety of checks in an attempt to verify the suspected cause of the failure. These checks include visual scan, infrared scope scan, individual die re-test, and DC measurements, as well as mechanical inspections of various kinds. When repairs are completed, the module is generally returned to Module Test for re-evaluation. If the failures were fixed, the open item is marked as resolved in the module log book. If repairs were

unsuccessful, or new failures were introduced, the appropriate entries are made in the open items list and the log book.

7.2.2 Full-Power Testing

In full-power testing the modules undergoing test must be installed into the liquid-cooled operating environment for which they are designed. Full-power testing allows the testing of all CRAY-3 module types. The minimum complement of modules which must be in the system for proper electrical and thermal loading consists of two Background Processors, one Foreground Processor, one I/O module and a full memory octant. The objectives of full power testing vary depending on which module type is the main subject of the testing. In all cases however, the overall goal of full-power testing is to provide performance data for the selection of modules to be installed into fully functional systems. Additionally, design issues and other engineering concerns can be addressed in the "real world" environment only with full-power testing.

Full-power testing requires a tremendous increase in the amount of equipment required compared to pulse-power testing. All of this increase is related to providing the liquid-cooled environment and continuous power for multiple modules. The actual quantity of diagnostic equipment decreases since nearly all testing is done with the use of diagnostic software rather than oscilloscopes, probes and digital multimeters.

TABLE 18. Equipment Used in Full-Power Testing.

Equipment	Function In Testing Procedures
System Cabinet.	The tank provides the liquid-cooled environment for the modules. The linear power supplies and translator cards are located below the modules.
The Control Pod.	Contains the coolant reservoir, pumps, and chilled water heat exchanger. The C-Pod also contains the Frequency Generator which serves as the system clock, the environmental control panel, and monitoring circuitry for electrical and thermal operating parameters.
Macintosh Computer.	Serves as the system console. It is connected to the Foreground Processor via cabling and a special interface card. All test software is loaded into the system through this console.
Oscilloscope.	An oscilloscope may be used to observe signals on an operating CRAY-3 system by employing inductive probes on the wireharnesses. The probes must be connected when the system is dry and the acrylic cover removed. Connections from the probes to the oscilloscope are made via one of the tank bulkheads.

Equipment	Function in Testing Procedures
Digital Multimeter.	The DMM is used primarily for checking continuity on signal paths and to measure transmitter-receiver resistance. The DMM can only be used when the system is shut down, the upper tank pumped dry, and the acrylic lid removed.
Hand Tools.	Adjustable torque screw driver. Connect and disconnect tools for the module wireharness connectors. Inductive probes. Tools required for attachment and removal of the acrylic tank lid, power buses and other hardware.

Before full-power testing may commence, the two processors, the Foreground Processor, and the I/O module are installed on the logic power bus. The 32 memory modules are installed on the memory power bus. Wireharnesses appropriate for the particular testing to be done are connected between the Background Processor modules, Foreground Processor module, I/O module and Common Memory modules. The acrylic lid is then placed on the top plate. It is fastened against the top gasket with non-metallic bolts. The liquid coolant is pumped into the tank from the C-pod filling the pedestal and the acrylic lid to a specified level. The circulation pump is then enabled. A pressure differential develops between the area confined by the acrylic lid and top plate, and the pedestal below, ensuring a positive flow of coolant through the system modules and causing all air bubbles to be expelled from the lid. Once the lid is completely full of liquid, power can be applied to the modules through the C-pod control panel. If no faults are detected, the power supply voltages and other environmental parameters will stabilize within a matter of seconds. The system is now fully powered and testing may commence using the Macintosh console. As in pulse-power testing, the System Test Interface application is run and the desired application program selected from the STI menu. Figure 128 provides a functional diagram of the software associated with the System Test Interface.

FIGURE 128. Block Diagram of the System Test Interface.




Table 19 lists additional software used in full-power testing which has not already been discussed in the preceding section on pulse-power testing.

TABLE 19. Additional Software Used In Full-Power Testing.

Program Name	Functions Performed
Background Macintosh Assembler.	The BMA allows a user to write and assemble off-line diagnostics for the Background Processor. It is used in conjunction with BMM.
Background Macintosh Monitor.	The BMM is a dynamic monitoring system that manipulates and observes the Background Processor functions. This mechanism allows a user to dynamically run CMT, BacPac, and BMM Interface in multiple processors.
Foreground Macintosh Assembler.	FMA enables a user to write and assemble off-line diagnostics for the Foreground Processor. It is used in conjunction with FMM.
Foreground Macintosh Monitor.	The FMM is a monitor specifically designed for use in the Foreground Processor. It allows the user of the diagnostic to control when the diagnostic test starts and how it operates.
Basic Address Adder Test.	The BAAT test performs a very basic check-out of an address adder.
Basic Address and Scalar Register Test.	BAST uses a floating ones pattern to check the validity of the address and scalar registers.
Basic Branch Addressing Test.	BBRT confirms the ability to jump to and from all locations in memory from all registers. It also checks the validity of all jump instructions from all registers.
Basic Common Memory Test.	BCMT performs a basic Common Memory test and provides error logging.
Basic Floating Point Test.	BFPT performs a basic test of the floating point functional units. Floating point instructions are tested against a set of pre-determined operands and results.
Basic Local Memory Test.	The BLMT program tests background local memory using only scalar reads and writes.
Basic Memory Vector Test.	BMVT tests the vector register's 16 x 4's with patterns from memory.
Basic Scalar Logical Test.	BSLT uses a floating ones pattern to check validity of the scalar logical unit.
Command Simulation Test.	CMD is designed for use as a CPU confidence test, a preventative maintenance tool and for use in isolating intermittent failures, particularly multiple instruction sequence problems.

Program Name	Functions Performed
Comprehensive Milner Algorithm Test.	CMIL uses a Milner algorithm to perform Common Memory integrity testing.
Comprehensive Register Conflict Test.	CRCT is a random instruction register conflict diagnostic which checks register conflicts/control problems.
Comprehensive Random Instruction Test.	CRIT
Comprehensive Slow/Fast Conflict Test.	CSFT is a random instruction (fast/slow) diagnostic which checks register conflicts/control problems.
Comprehensive Vector Instruction Test.	CVIT is a vector register test that uses different path, first pass and a bit counter method to detect register and functional unit failures.
Extensive Address Multiply Test.	EAMT allows for extensive testing of address multiply.
Extensive Address and Scalar Test.	EAST allows for exhaustive testing of the address and scalar registers.
Extensive Memory Backup Test.	EBUT tests Common Memory interface back-up buffers using single and dual-port vector commands.
Extensive Common Memory Test.	ECMT is an extensive Common Memory test which correlates to foreground based FMM diagnostics FCMT and FCCM. Banks are tested in parallel and patterns are targeted to individual chips.
Extensive Floating Point Test.	EFPT simulates all results to verify functional units.
Extensive Instruction Buffer Test.	EINT
Extensive Local Memory Test.	ELMT tests background local memory using scalar and vector read/write.
Extensive RA Table Test.	ERAT
Extensive Scalar Adder/Pop/LeadingZ Test.	ESAT extensively tests the scalar integer adder functional unit.
Extensive Scalar Shift Test.	ESHT is used to exhaustively test all scalar shift instructions.
Extensive Simulate Result Test.	ESMT tests all A and S register non-floating point and VM.
Extensive Semaphore and Status Register Test.	ESST tests the status register, limit register, EXIT instructions, 004 through 007 instructions and 035 instructions.
Extensive Vector Scalar Comparison.	ESVT performs a random compare of vector versus scalar instructions to check for compatibility of answers.
Extensive Vector Adder/Iota Test.	EVAT tests the vector integer adder functional unit as well as the compress iota.

Program Name	Functions Performed
Extensive Vector Shift/Pop/LeadingZ Test.	EVBT tests all vector shift instructions as well as the vector population/leading zero units.
Extensive Vector Logical Test.	EVLTL thoroughly tests the logical section of the vector functional unit. All vector mask and vector logical instructions are tested. A RTC test is also included.
Extensive Vector Memory Test.	EVMT tests the background Common Memory using vectors.
Extensive Vector Register Module Test.	EVRT is a random data vector register test.
Foreground Adder Test.	FADD performs a basic checkout of the foreground adder functional unit.
Foreground Address Register Test.	FBAR tests the Background Processor A registers via the Foreground Processor and channel loops.
Foreground Scalar Register Test.	FBSR tests the Background Processor S registers via the Foreground Processor and channel loops.
Foreground Clears Local Memory.	FCLR is a utility test program used to verify FMM operations and clear local memory.
Foreground Constant Data Test.	FCON tests constant instructions 20, 21, 22, 30, 31, 32, 200, 220, 300, 320 and logical instructions 44,45 and 46.
Foreground Instruction Timing Test.	FCONFLICT tests conflict timing conditions in the Foreground Processor.
Foreground Conflict Integrity Test.	FCRPT
Foreground Background Communication Test.	FGBC tests the interrupt circuitry by loading and running short programs in the Background Processor which create error conditions. This includes testing of the base, program and limit addresses.
Foreground Channel Function Test.	FGCF tests the I/O logic that forms the interfaces between the Foreground Processor, external devices, Background Processors and Common Memory.
Foreground Common Memory Test.	FGCM tests Common Memory via the Foreground Processor and channel loops.
Foreground Status and Response Register Test.	FGDH tests the DH stack status register, the call pointer address and all the interrupt conditions from the Foreground Processor.
Foreground Local Memory Test.	FLAT
Foreground Logical Test.	FLOG performs a basic checkout of the foreground logical functional unit.
Foreground Register Test.	FREG performs a basic A, B, C and conditional jump test. No functional units are used in the test.

Program Name	Functions Performed
Foreground Real Time Clock Test.	FRTC tests the real time clock by checking consecutive clock cycles at random intervals, using incremental delays, and testing for stuck at zero and stuck at one conditions for all 32 bits.
Foreground Shifter Test.	FSHF tests the foreground shift functional unit.
Foreground Instruction Buffer Test.	FXIA tests the instruction buffers from the Foreground Processor and the exit jk path.

The software tests listed in Table 19 are all off-line diagnostics; that is they are run by themselves, while the operating system is not running. As was the case with pulse-power testing, errors are detected when the expected data for a given test packet differs from the data received from the module under test. The application will report the error on the Macintosh screen. Errors may be logged and stored to a report file using the Failure Consolidation Package.

Modules on which failures have been identified in full-power testing are generally sent to Module Test for verification and troubleshooting under pulse power. The module will then remain in the pulse-power test/repair loop until these open items are resolved. At that time they will be returned to full-power testing. However, there are times when the diagnosis of a problem can only be properly implemented in full-power testing. Under these circumstances a module may go directly to repair and return directly into the full-power environment, completely bypassing the pulse-power test loop.

7.3 Module Repair

The electronic section of the CRAY-3 is designed to be repairable down to the individual die in a module assembly. However, for installed systems the unit of exchange will normally be a block of modules. This replacement procedure will eliminate the time that would be needed to disconnect and then re-connect the wiring harness if a single module was to be removed for repair.

As we have seen, the software diagnostics for the CRAY-3 analyze a fault down to the individual die or group of die. Copies of the analysis documentation are sent along with the module to be repaired. This documentation informs the people who will be performing the repair of the known or suspected cause of failure. Based on that documentation, along with the repair history for that module, a decision is made as to the best approach for each repair. Sometimes a entire stack will be replaced even though only one die is suspected to be faulty. This may be done, for example, because since this

module was built improvements have been made to several other die on that particular stack.

With an individual module assembly, the unit of stacking and unstacking is one single board stack. The module is placed in an assembly pallet and the guide, logic and power pins associated with the stack needing repair are removed and discarded. The four circuit boards in the stack are now separated and made available for examination. If the problem is a faulty die, that die is removed. This is accomplished by attaching a steel fixture with adhesive to the back of the faulty die which is then pulled from the printed circuit board. A new die is now inserted. The normal die attach process is used with modifications in the die attach fixture to accommodate the 15 die which are already attached to the circuit board. On completion of the die attach process the stack is re-assembled on the module using the Universal Stacking fixture. The module is now re-submitted to pinning so that new logic jumper and new power pins may be inserted. The repair cycle time for a single stack is less than a normal working shift. The module is now available for Module Test or for re-submission to a system.

7.4 Failure Analysis

During the module repair process, components are inspected to determine the cause of the reported failure. This information is logged and analyzed so that trends can be spotted that might give some clue as to how the manufacturing process or the design of the component itself could be improved.

The failure analysis procedure starts with a list of failing components along with a description of the problem found in Module Test. This information is entered into a Macintosh-based database tracking system. This module tracking system generates forms that are filled out at each stage of the repair process.

As a module is unpinned and then unstacked, each board is inspected for damage, contamination, or manufacturing errors. Defects are noted on the forms that travel with the module. Failures that involve die are sent to Die Attach for inspection.

Before a die is removed it is inspected with an infrared microscope to check for damage to the substrate or bonding pads. After the die is pulled off the circuit board, the gold leads are inspected to make sure the bond is intact and that the stamping forces and stamping procedures were correct. Finally, the micro-circuits of the die are inspected with a high-power microscope to check for scratches or fabrication defects in the circuits.

When this process is completed, a two-letter failure analysis code is assigned to each component identified for repair. These codes are entered into the same database tracking system that generated the repair travelers which accompanied the module. This database containing the history of every repair can then be searched and sorted by different criteria and appropriate reports generated. By studying failure modes and trends observed in the repair process changes can be implemented to improve the manufacturing of new modules. The repair codes which are used are listed in Table 20.

TABLE 20. Failure Analysis Codes.

Code	Code Name	Description
BF	Bond Failure.	Bond failure due to poor welding.
BI	Bad Insert.	Poor seating height, bad cut, etc.
BP	Bad Pin.	No bulge or some other problem with a pin.
BW	Broken Wire.	Wire broken at printed circuit board or connector.
CB	Cratered Bond.	Cratered bond pad discovered at IR inspection.
CR	Cracked.	Cracked or chipped die.
DA	Die Attach.	Die attach mechanical failures; shear, mis-stamp, etc.
DM	Debris in Module.	Metal contamination found in module at unstack.
EF	Electrical Failure.	Suspected electrical failure. Requires FAST re-testing of suspected die.
IN	Inconclusive.	Die evaluated, but no damage was found.
MP	Missing Pin.	No pin was inserted during Module Pinning.
NA	Not Analyzed.	Die was not analyzed. No data is available.
NF	No Failure.	No problems found during inspection.
OL	Open Lead.	Pad lift or bad solder connection.
OM	Other Module Assembly.	All other failures found at Module Assembly.
OP	Open.	Open connection. Die attach mechanical failure, pad lift, or cratered bond is suspected.
OS	Other Solder.	All other failures found at Solder.
PC	Printed Circuit Board.	Problem with PCB thickness, warpage, hole diameter, drilling, damage or circuit trace short or open.
PL	Pad Lift.	Pad lifting discovered at IR inspection.

Code	Code Name	Description
SC	Scratch.	Scratch on circuit traces. Metal has been damaged.
SE	Stacking Error.	Boards swapped or rotated. Missing spacer or flowblocker.
SL	Short in Logic Plate.	Short or open in logic plate.
SO	Short in Ohmega Plate.	Short or open in resistor plate.
SP	Short in Power Plate.	Short or open in power plate.
TP	Twisted Pair.	Twisted pair inversion.
UC	Upside Down Connector.	Connector installed incorrectly.
UP	Upgrade.	Upgrading die, plates or boards for new revisions or for new testing.

We have now completed our discussions of the major components that make up the CRAY-3 electronic modules. We have also presented a brief discussion of how these assembled modules are tested and repaired. Having done that we can move on to Chapter 8 where we introduce the much larger pieces of the machine that hold all of the smaller pieces and allow them to function properly.

Cabinet Design

In this chapter we turn our attention away from the smaller component parts of the CRAY-3 and look at the electrical and mechanical systems that tie everything together to make it a functional machine. We will discuss the System Cabinet and its associated components, including the Octant Assembly, Wire Harnesses, Cooling System, Translator Cards and Power Supplies. After those discussions we will look at the System Control Pod and discuss the functions which it performs in cooling, controlling and monitoring the entire machine. Finally, we will present an overview of peripheral equipment which is typically used with the CRAY-3.


8.1 System Cabinet

The CRAY-3 computer system is contained in a series of modular tanks arranged in varying configurations. A front view of a single tank is illustrated in Figure 129. The lower section of the tank contains the power supplies for the CRAY-3 circuits and the cables for communication with peripheral equipment. The section above the tank top plate contains the electronic modules for the computer.

The System Cabinet external footprint for a four-processor version is 109.22 cm wide (43 inches) and 71.12 cm in depth (28 inches). The System Cabinet external footprint for a 16-processor is a 109.22 cm (43 inch) wide octagon. The power supply section stands 100.33 cm (39.5 inches) above the false floor.

The acrylic lid extends the height another 21.59 cm (8.5 inches) giving a total height above the false floor of 121.92 cm (48 inches).

FIGURE 129. Front View of a Four-Processor Tank.



The power supply portion of the cabinets employs a two-piece cosmetic skin which can be separated and removed for access. The electronic section at the top of the cabinet is covered by a one-piece acrylic cover. The cover is attached to the tank top plate with non-metallic screws. The acrylic top is sealed to the tank top plate with an integral O-ring. The tank assembly is designed to stand on the sub-floor, below the false computer room floor. The legs on the baseplate are designed to allow the system to adjust to various floor heights so the system skins fit neatly on the false floor tiling (see Figures 133 and 134).

FIGURE 130. Front View of Eight-Processor System Cabinet.





FIGURE 131. Eight-Processor System Cabinet with Skins.



The section under the computer room floor is organized into bulkheads for cable and cooling connections. Above the computer room floor, the power supply section and the electronic module section are organized into individual sections called octants (tanks). Each system cabinet has from one to eight octants depending on the system configuration. Each octant is in effect a one-eighth slice of an octagonal pie. Each tank is a single-piece aluminum casting. The tanks are fastened to the baseplate which serves as a manifold to evenly distribute coolant to each tank. The area enclosed by the acrylic lid, above the top plate, is common to all of the octants.

FIGURE 132. Front View of Prototype One- or Two-Processor System Cabinet.





FIGURE 133. Sub-Floor Mounting of the System Cabinet.



As mentioned earlier, the construction of the system is modular. Individual octants can be added to the baseplate to form groups of one, two, four and eight tanks (see Figure 134). In this way an initial configuration of a single processor (in one octant) can be upgraded to a 16-processor machine (in eight octants). The electronic assembly section above the top plate is also arranged into octants (see Figures 135, 136, 137 and 138).

FIGURE 134. Baseplate for the Octants.

The octants, which hold the power supplies and modules, fit on the baseplate. The baseplate shown here can accommodate up to four octants (eight processors with their associated memory and power supplies).

One octant would fit in this area.




FIGURE 135. Top Plate for a Four-Processor CRAY-3.




FIGURE 136. Top Plate for an Eight-Processor CRAY-3.




FIGURE 137. Modules Mounted In an Eight-Processor Cabinet.






FIGURE 138. Top Plate for a Prototype One- or Two-Processor CRAY-3.



8.2 Octant Assembly

As we have seen there are 336 modules in the upper portion of the CRAY-3 cabinet. These are grouped in eight octants of 42 modules each. Each octant has two ranks of modules. The outer rank of modules is for Common Memory and consists of 32 modules. The inner rank of 10 modules is for the Background Processors, I/O modules and one Foreground Processor per machine. Each octant contains two processors of four modules each and has provision for two Foreground Processor or I/O modules. The individual modules are mounted vertically into a horizontal power bar and are spaced on 7.62 mm centers. The modules are interconnected by means of connectors attached to twisted pair wires or Flex circuits soldered to the outer boards of each module. These 44-pin logic connectors are then interconnected by means of a system wire harness. There are local connections direct from module to module and also interconnections from octant to octant by means of the system wire harness.

FIGURE 139. Modules and Power Supplies In an Octant.



Electrical power is delivered from the power supply section to the electronic module section via vertical power bus bars. Horizontal power buses attached to the vertical power buses coming from the power supplies then accept the four tapered power blades of the modules allowing for both a mechanical and electrical connection to the buses. The four power blades on each module are forced into slots in the power distribution blocks by screws. The power distribution in an individual octant is illustrated in Figures 140, 141, 142 and 143.

FIGURE 140. Power Distribution in an Octant Assembly.




FIGURE 141. Horizontal Power Buses for Memory from Module Side.




FIGURE 142. Horizontal Power Buses for the Memory Modules.




FIGURE 143. Horizontal Power Buses for Logic Modules from Module Side.





FIGURE 144. Horizontal Power Buses for Logic Modules.



The octant assembly of modules is removed from the cabinet as a complete section. The outer rank of memory modules with the horizontal power blocks is removed as one unit after disconnection from the system wire harness. The center block of processor and I/O modules is also removed as one group in the same way as the memory modules.

In the next section we will look at the power supplies that provide the necessary voltages for the modules at the current demands required for continuous operation of the computer.


8.3 Power Supplies

The CRAY-3 computer system uses a bank of eighteen linear power supplies to provide power to the memory and logic modules for each octant. Each power supply provides approximately 1,700 watts at either +3.3 volts DC (known as Vp), -1.2 volts DC (known as Vn) or +5 volts DC (known as Vm).

The power supplies are grouped into six racks to simplify power distribution and ease of installation. Each of the two Vp racks used for powering the logic modules holds four power supplies. The two Vp racks used to power the memory modules hold seven power supplies. One rack is used to support a single Vm power supply. The sixth rack is used to hold two Vn power supplies. This sixth rack is physically split to allow the output of one Vn power supply to provide Vn to the logic modules and the other Vn power supply to provide Vn to the memory modules. These racks of power supplies are located in the System Cabinet directly below the modules. The power supplies are immersed in the coolant that passes through the modules prior to its return to the Control Pod heat exchanger.

The power supply output power is connected to the modules via a system of vertical and horizontal bus bars (see Figures 140, 141, 142, 143 and 144). The input power is provided to the power supply racks via flexible, ribbon connectors that pass through a sealed penetrator and connect to a rigid vertical bus bar located behind the racks. The input power used by the power supplies is nominally rated at 220 volts AC, three-phase, 400 Hz. It is provided by an external motor generator (MG) and powers a wye and delta transformer in each power supply, resulting in twelve output phases from the transformers. Each of these phases is then rectified and their outputs combined. This voltage then passes through an output choke and is ready for use. The front plate of the power supply rack is used to parallel all input voltages for the supplies in that rack.

FIGURE 145. Linear Power Supplies In an Octant Assembly.



Since the power supplies are linear, output adjustment is accomplished by varying the input voltage. This is done in one of two ways. For Vn and Vm a voltage vernier board is used. This board uses an external triple ganged potentiometer to adjust the base drive on six bi-polar pass transistors. The action of these transistors will vary the amplitude of the input voltage passed to the power supplies. Two of these boards are used in the Vn rack and one in the Vm rack. Due to the much higher output current consumption of the Vp racks the vernier board is not used. Instead, the motor generator output voltage is adjusted through the use of its own controls. Since the power supply input voltage is common to all racks, the adjustment of Vp will cause a change in both Vm and Vn. Therefore, Vp is adjusted prior to setting Vn and Vm.

FIGURE 146. Simplified Top View of a Vp Logic Linear Power Supply.




FIGURE 147. Simplified Top View of a Vp Memory Power Supply.




FIGURE 148. Simplified Top View of a Vm Power Supply.





FIGURE 149. Simplified Top View of a Vn Power Supply.



Next we look at the wire harnesses used to interconnect modules, the processors and memory so that each module can perform the necessary logical functions as part of a complete computing system. The wire harnesses carry data and control signals. They do not deliver power voltages to the modules. As we have seen in this section power delivery to the modules is accomplished by the rack mounted linear power supplies and their associated system of fairly large metal bus bars.

8.4 Wire Harnesses


In a CRAY-3 system each of the octants has nine unique wire harnesses electrically connecting the modules to each other. These nine wire harnesses are grouped in three functional sets: the Common Memory, the first Processor (P0) of the octant and the second Processor (P1) of the octant. Each functional set is comprised of a top, center (they connect to the edge opposite the power blades) and bottom wire harness.

The top wire harness for Common Memory is comprised of 144 connectors which mate with the modules. There are approximately 1,270 connections in this set. The center wire harness for Common Memory is comprised of 507 connectors which mate with the modules and 62 connectors referred to as "floating connectors." There are approximately 4,830 connections in this set. The floating connectors serve as links between Common Memory, the Processors and other octants. The bottom wire harness for Common Memory is comprised of 154 connectors which mate with the modules and eight floating connectors. There are approximately 1,390 connections in this set.

The first and second processors of the octant are connected with wire harnesses on the top, center, and bottom of the modules. Each processor has unique wire harnesses in the octant. There are approximately sixty percent fewer connectors in these wire harnesses than in the Common Memory harnesses. This accounts for approximately fifty percent fewer total connections in these wire harnesses. As additional octants (hence processors) are added to the system, the center wire harnesses are altered to accommodate floating connectors for communication between octants.

Each module has a specific number of logic connectors which allow signals to communicate with other modules in the system. The number of edge connectors on a module ranges from 24 on the F and H modules (Common Memory) to 46 on the B module (part of a Background Processor). These connectors are mated with a wire harness to allow communication with other modules. Each grouping of modules has wire harnesses that connect to the top, center and bottom connectors on each module. This is illustrated in Figure 150.

FIGURE 150. Wire Harness Connection Points to the Modules.




The connectors used are a molded plastic part with two rows of 22 pins. One row of pins connects to the true (logical zero in the CRAY-3) wires of a twisted pair of 34 AWG wire. The other row connects to the false (logical one in a CRAY-3) wires of the twisted pair. There are 22 twisted pairs of 34 AWG wire on each connector.

The production process for each wire harness begins by stripping and pre-tinning each twisted-pair on the connector to the proper lengths specified by the design documentation for that wire harness. Each connector (with the corresponding twisted pairs of wires attached) is then loaded into a fixture

which allows them to be oriented in the same relationship to each other that they will have when attached to the modules in an octant. Each wire is located using an electrical homing device and connected to its signal destination using a solder sleeve.

The solder sleeve is a piece of translucent shrink tubing approximately 5 mm long with a ring of solder and flux. When heat is applied the solder is re-flowed making a solder connection as the tubing shrinks tightly around the connected wires. Each connection is tested to ensure its quality.

FIGURE 151. Simplified Diagram of Four-Processor Wire Harness.



Note: Curved, bold lines represent bundles of wires in this illustration.

8.5 Cooling

The CRAY-3 is cooled by direct contact immersion in an inert liquid fluorocarbon similar to that employed in the CRAY-2. However, the high power density of the CRAY-3 required further development of the CRAY-2 cooling mechanism to provide adequate heat removal. Operating temperature throughout the computer circuits averages 30° Celsius.


The high packaging density of the CRAY-3 results in a power density of up to 640 watts per cubic inch. Each nine-layer CRAY-3 logic module is 121 by 107 by 7.14 millimeters and contains up to 1,024 gallium arsenide logic circuits consuming from one to three watts each. Total power consumption for a two-processor octant is approximately 45,000 watts.

The individual modules are separated from each other by insulating separators or shims. These are sheets of acrylic material slightly larger than the module face and 0.203 mm thick. These shims block any inter-module coolant flow so that all coolant flow will be directed between the layers of each module rather than between the modules themselves. This flow blocking insures that the coolant will flow over the die and around the exposed areas of the jumpers for maximum heat transfer.

A key factor in the design of the cooling mechanism is the narrow channels between the module layers through which the cooling fluid must flow. The channels are a maximum of 300 microns from the back surface of a die to the adjacent printed circuit board. The flow through the channels is also partially blocked by many z-axis interconnect twist pins distributed throughout the modules. The pins actually serve to improve the heat transfer ability of the fluid by continually breaking up the insulating boundary layer that would typically form with flow between two flat planes.

The cooling system utilizes a closed loop re-circulation design. Fluid cooling and conditioning is performed in the C-Pod, away from the System Cabinet. The operation is best described by tracing the fluid path through the primary components in the loop (refer to Figure 152). The circulation pump moves warm fluid returning from the tank through a heat exchanger and through conditioning filters in the C-Pod. The cool fluid is then delivered to the computer tank baseplate via piping under the computer room false floor. The fluid is routed through the baseplate and ported into the area above the CRAY-3 electronic modules and forced down through the modules. As the fluid exits the bottom of the modules into the top of the tank the cooling of the computer circuits is complete. The fluid then flows down through the tank cooling the system power supplies prior to exiting the bottom of the baseplate. The warm fluid is returned to the pump inlet in the C-Pod to complete the circulation loop. Total temperature rise of the coolant across the loop is approximately 5° Celsius.

FIGURE 152. Simplified Diagram of the Cooling System.



A fluid reservoir is located in the C-Pod to serve as a holding tank while servicing one of the tanks. When the system is not operating, any one or two tanks may be simultaneously emptied for servicing. The reservoir is also used to supply continual make-up fluid to the circulating loop during system operation.

FIGURE 153. C-Pod Cooling Components—Front Right View.





FIGURE 154. C-Pod Cooling Components—Front Left View.




8.6 Translator Cards

The translator cards are used to convert the signal levels used in the CRAY-3 into signal levels required by any peripheral equipment connected to the machine. The translator cards also convert signal levels from the peripheral devices into signal levels required by the CRAY-3. All I/O signals to and from the CRAY-3 pass through the translator cards with the exception of the console interface where the CRAY-3 signals are directly connected to the console equipment via a custom interface board.

Presently, the translator cards and their associated power supplies are mounted in an electronic components box located underneath the computer room tile floor (see Figure 155) within six feet of the System Cabinet. The translator box mounts on a small stand which in turn sits on the sub-floor in the computer room.

Each translator card is structured to provide a maximum of 48 input/output signal pairs. These 48 pairs are divided into two groups in the documentation to provide clarity. There are three different types of translator cards: low-speed, high-speed HIPPI, and high-speed. The low-speed cards are used for the I/O signals passing to the DD-49 or LSX equipment. Two DD-49 or LSX Systems can be interfaced to a single translator card. The high-speed HIPPI cards are used when connecting HIPPI devices. A single translator card is required for each HIPPI device. The basic difference between the low-speed and high-speed HIPPI cards is in the Emitter Coupled Logic (ECL) termination used. The low-speed cards use a 60 ohm to -2 volt termination whereas the high-speed card terminates the ECL signals in accordance with the HIPPI specification. The high-speed HIPPI cards also differ from the high-speed cards in the way specific HIPPI control lines are terminated in the CRAY-3 interface implementation. The HIPPI translator card has been modified so that the interconnect signal to the HIPPI device is pulled to a ECL low. This was done for both the source and the destination cables.

FIGURE 155. Translator Box.



8.7 System Control Pod

The ancillary functions for the CRAY-3 are housed in a separate cabinet known as the System Control Pod. The System Cabinet containing all of the logic modules and their associated power supplies does not have any control capabilities. Therefore, all of the monitoring and control must be done through the C-Pod and the System Console. The C-Pod measures about 127 cm (50 inches) square by 165.1 cm (65 inches) high and resides in the same room as the CRAY-3 System Cabinet. Approximately one-fourth of the System Control Pod sits below the raised computer room floor leaving about 138.4 cm (54.5 inches) visible in the computer room.

The System Control Pod is connected to the main CRAY-3 system cabinet under the computer room floor by means of power and control cables, and coolant piping. There are three cables per octant fitted to the CRAY-3 cabinet. The fluid transfer for coolant between the CRAY-3 cabinet and the Control Pod is by means of two 3-inch hoses; one for fluid flow into the CRAY-3 cabinet and one for fluid flow back from the cabinet.

The System Control Pod houses electronic controls and monitoring devices for both the electrical and fluid handling equipment of the CRAY-3. The logical operations of the CRAY-3 system and system Dead Start functions are monitored and performed from the System Console. The console is a microprocessor-based UNIX workstation. The workstation is connected to the Foreground Processor of the CRAY-3 through the translator cards located in the translator box underneath the computer room floor.

The electronic equipment in the System Control Pod includes the front panel, the maintenance panel, the electronic module, the motor starters, the clock source, the AC box and the DC bulkheads. The components of the fluid handling system include the reservoir, the heat exchanger, the up/down pump, the circulate pump and the filters. The reservoir stores the Fluorinert when it is not in the tank. Initially the tank is filled with Fluorinert from the reservoir via the Up/Down Pump. After the tank is full and the control valves for the Up/Down Pump are closed the Fluorinert can be circulated through the system by the Circulation Pump. The Heat Exchanger continually cools the Fluorinert as it circulates through the system. There are two filters in the system. One removes any particulates or debris while the other handles chemical by-products that will be present if the Fluorinert reaches an elevated temperature and burns. The following section will only address the electronic equipment in the System Control Pod.

FIGURE 156. C-Pod External View.






FIGURE 157. C-Pod-Front Right View.



Note: C-Pod Components are identified in Figures 153, 154, 163 and 164.

FIGURE 158. C-Pod Front Left View.



Note: C-Pod Components are identified in Figures 153, 154, 163 and 164.

8.7.1 Front Panel User Controls

The front panel of the System Control Pod is the primary means by which the user will operate and control the non-logical operations of the CRAY-3. The front panel can be broken down into three sections (see Figure 160). The top portion of the front panel contains the touch screen display panel. The center section of the front panel contains the user switches board which holds the main switches for operation of the CRAY-3. The bottom section of the front panel contains the maintenance switches board which holds the auxiliary switches that are used less frequently.

FIGURE 159. Positioning of User Controls and Maintenance Panel.





FIGURE 160. C-Pod User Controls.



There are six switches on the front panel for primary operation of the CRAY-3. The switches are mounted to the user switches board which is located behind the front panel. These six user switches include 1) the on/off switch for the circulation pump, 2) the on/off switch for the motor generator power, 3) the on/off switch for the motor generator exciter, 4) the direction of the up/down pump, 4) the system reset switch, and 6) the lift switch for the top cover.

There are eight switches located at the bottom of the front panel for secondary operation of the CRAY-3. These switches are mounted on the maintenance switches board that is located behind the front panel. The eight maintenance control switches (for up to an eight-processor system) operate 1) motor starter reset, 2) the pressure interlock, 3) voltage sense, 4) auto-leveling, 5) octant valve position for Octant 0, 6) octant valve position for Octant 1, 7) octant valve position for Octant 2 and 8) octant valve position for Octant 3.

The operating parameters which are monitored for the CRAY-3 are displayed on the touch screen display panel located at the top of the front panel (see Figure 161). This includes the conditions of the fluid handling equipment (flow meters, pressure transducers, thermocouples, level sensors and the status of the pumps), the output voltage for the motor generator and the DC bus voltages in the System Cabinet.

FIGURE 161. The Touch Screen Display.

SYSTEM		TOP PRESSURE : 3.6 PSI DIFF PRESS OCT0 : 3.2 PSI DIFF PRESS OCT1 : 3.0 PSI FC FLOW : 170.5 GPM HxH2O FLOW : 175.6 GPM				CIRC INLET PT : -1.6 PSI CIRC OUTLET PT : 83.2 PSI CIRC INLET TC : 19.9 C	
MG VOLTAGE	: 216.7	HxFC OUTLET PT	: 60.0 PSI	HxH2O OUTLET PT	: 18.8 PSI		
MG CONTROL	: ON	HxFC INLET TC	: 22.0 C	HxH2O INLET PT	: 0.0 PSI		
MG EXCITE	: ON	HxFC OUTLET TC	: 16.0 C	HxH2O INLET TC	: 12.1 C		
MG READY	: YES						
		OCT0	OCT1				
VP MEMORY	3.30	3.32	CIRC PUMP	: ON	RES HIGH	: DRY	
VN MEMORY	1.21	1.20	PU/PD PUMP	: OFF	RES WARN	: DRY	
VM MEMORY	5.01	5.05	VOLTAGE SENSE	: ENABLE	RES LOW	: WET	
VP LOGIC	3.30	3.42	UpDnPUMP TRIP	: OK	TANK HIGH	: DRY	
VN LOGIC	1.20	1.21	CIRCPUMP TRIP	: OK	TANK WARN	: WET	
					TANK LOW	: WET	
					PD OCT0	: WET	
					PD OCT1	: WET	
H2O FLOW LOW							

The flow meters measure the rate of flow of both the chilled water and the Fluorinert in gallons per minute. The thermocouples measure the temperature, in degrees Celsius, of the Fluorinert and the chilled water at various points throughout the system. The pressure transducers measure the pressure in


pounds per square inch in the tank, the heat exchanger, and at the pumps. The level sensors measure how much fluid is left in the tank or the reservoir. The status of the pumps (powered or tripped) is also recorded.

Both AC and DC voltages in the system are monitored and displayed on the front panel of the C-Pod. The motor generator status is also displayed along with the value of the output voltage (phase-to-phase) of the motor generator. The values of the DC bus voltages from the power supplies in the tank are also monitored and displayed.

8.7.2 Maintenance Panel

The maintenance panel of the Control Pod is located to the left of the front panel and contains several controls for adjusting operating parameters of the CRAY-3. The output voltage of the motor generator, the flow rate of the Fluorinert, and the DC bus voltages in the tank (except for V_p) can be fine tuned by the user from the maintenance panel. The knobs that control the DC bus voltages in the tank are mounted on the voltage adjustment boards. There is one voltage adjustment board for each octant and they are located just behind the maintenance panel. The user may adjust V_p by adjusting the motor generator output until the desired V_p is attained. Once the V_p voltage is set the user may adjust the voltages for V_n Logic, V_n Memory, and V_m Memory. The maintenance panel also contains a clock that tracks the amount of time the system has been operating.

FIGURE 162. Maintenance Panel Controls.



8.7.3 Electronic Module


The electronic module is located behind and to the left of the maintenance panel. It contains three sections (see Figure 163). The top section contains the power supply plate. The middle section contains the card cage for the PC boards and the bottom section consists of three signal condition boards.

The power supply plate consists of a multiple output power supply and the power supply board. The power supply provides +5 volts, +12 volts and -12 volts DC. Power is supplied to the mother board, the touch screen display board, the top cover lift, the DC bulkhead, and the signal condition boards in the System Control Pod.

The card cage, located in the center of the electronic module, contains the control board, the display controller board, the filter board and the mother board. The control board monitors and controls operation of the fluid handling system. The display controller board monitors system parameters, displays system information and acts as an interface between the Control Pod and the system console. The filter board is used to filter signals from the tank, such as DC bus voltages, pressure transducers and other input devices. The mother board allows connection of the control board, the display controller board and the filter board with the rest of the electronics in the Control Pod.

There are three signal condition boards located at the bottom of the electronic module. The signal condition boards convert and condition the thermocouple signal and the signals of the flow meters and pressure transducers to voltage levels that the control system can use. Two of the boards are eight-channel boards and the third board is a four-channel board.

FIGURE 163. Electronic Assemblies In the C-Pod—Front Left View.



8.7.4 Motor Starters

There are two motor starters located at the bottom of the Control Pod on either side of the circulation pump. The motor starters are used to start the circulation pump and the up/down pump. The motor starter is actuated as follows: The switch on the front panel is connected to the control board. When the switch is activated and all conditions are acceptable, the control board sends a DC signal to the DC coil of the relay. The AC contact of the relay is connected to the motor starter which actuates the specified pump.

8.7.5 Clock Source

The system clock is supplied by the clock source located above the electronic module to the left of the front panel. The clock source is a Wavetek Model 2410 Synthesized Signal Generator. It is generally set at 500 mHz and 13 dBm. At the present time, the clock source may only be adjusted manually by the user. Future design changes will allow the clock source to be adjusted from the System Console (UNIX workstation).

8.7.6 AC Box

The AC power and the DC valve signals are transferred in and out of the System Control Pod via the AC box. The AC box is located below the raised computer room floor on the front side of the System Control Pod. The AC box is divided into two sections. The left section contains the power and signals for the valves, and the power control signal for the translator card assembly. The right section contains the power to power pumps and other equipment located in the System Control Pod.

The left section of the AC box contains the power relays, the C-Pod relay board and the AC Bulkhead. The relays are used to control the valves and starters. The C-Pod relay board acts as an interface between the control board and the valves and motor starters. The AC bulkhead transfers AC power out of the System Control Pod and DC valve signals into the System Control Pod.

The DC valve signals are transferred into the System Control Pod through the connectors mounted on the top section of the AC bulkhead. These 5 pin connectors for the DC valve signals are mounted to the valve signal board which is located behind the AC bulkhead panel. The DC signal connections are for valve positions of the octant's valves; namely, the up/down valve and the up/down bypass valve.

The AC power control lines are transferred out of the System Control Pod through the connectors mounted on the bottom section of the AC bulkhead.

These six pin connectors for the AC power control lines are hard-wired to the power relays. The power relays and the relay board control the position of the valves by sending power to the actuator for the valve. The AC power connections include power for the remote translator box, the octant valves, the up/down valve, the up/down bypass valve, and the compensation valve.

The right section contains the motor generator bulkhead board and terminal strips. There are no connectors on the right section. All power lines are passed through seal tight in the sides of the AC box and are hard-wired to the terminal strips. The 120 volts is then distributed to the relay board, the power relays, the clock source, and the power supply plate. The 480 volts is distributed to the motor starters which operate the circulation pump and the up/down pump. The motor generator bulkhead board acts as an interface between the control board and the motor generator, which is used to power the DC power supplies in the tank.

8.7.7 DC Bulkhead

The DC signals are transferred in and out of the Control Pod via connectors on the DC Bulkhead. This bulkhead is located below the raised computer room floor on the left side of the System Control Pod (see Figure 163). The connectors are mounted to the DC bulkhead board which is located behind the DC bulkhead panel. The signals on the DC connectors include the voltage sense from each octant, as well as the common level sense and top pressure from the tank. The Fluorinert and chilled water flow meter readings, the high voltage sense, the RS-485 communication signals and the RS-232 signals to the System Console are also transferred through the DC bulkhead.

FIGURE 164. Electronic Components in the C-Pod—Front Right View.





FIGURE 165. Sub-Floor Mounting of the C-Pod.



8.8 Peripheral Equipment

The equipment that will be discussed in this section includes data storage units which can be interfaced with the CRAY-3. Currently, three systems are supported. These are the DD-49 and DS-40 disk systems from Cray Research, Inc. and the RAID (Redundant Array of Inexpensive Disks) systems from IBM. A brief discussion of the CRAY-3 circuitry used to control and interface with these devices will be given following each hardware description.

8.8.1 DD-49 Disk Unit

The DD-49 Disk Storage Unit is manufactured by Cray Research, Inc. Each unit has a capacity of 9.75 gigabits of data. The DD-49 consists of nine rotating disks with 16 data surfaces. Data is accessed by 32 read/write heads organized into eight groups, four read/write heads per group. The average rotational latency is 8.3 milliseconds.

Sustained data transfer rate, using the four recording heads in parallel, is 78 megabits/second. Consecutive sectors of data can be read or written with alternating data buffers in the disk controller. The head group selection can be changed as data is moving for the last sector on the disk track. The first sector on the next track can then be processed without missing a disk revolution.

The heads are positioned by two identical head actuator (servo) mechanisms to one of 888 cylinders. The servo mechanisms are identified as Servo-A and Servo-B. All heads move together in the positioning process. Moving to an adjacent cylinder requires 2.5 milliseconds. Moving the maximum distance of data cylinders requires 30 milliseconds.

The recording surface available to each head group is a disk track, which is the basic storage unit implemented. Within each disk track there are 42 sectors in which data can be recorded and read back. The data in one sector is called a data block and consists of 512 64-bit words plus verification and error correction data. All data transfers are either 16 or 2048 parcels.

8.8.2 DS-40 Disk Subsystem

The DS-40 Disk Subsystem, also manufactured by Cray Research, Inc., consists of three components, the DCC-2 chassis, the DC-40 disk controller, and the DD-40 disk storage unit.

The DCC-2 chassis contains the interface logic to operate the DD-40 DSU. The chassis contains four dual ported DC-40 controllers that communicate with a Foreground Processor channel.

Each DC-40 controls two DD-40 disk storage units. Only one DSU controlled by a DC-40 will be active at one time. The DSU selection is controlled by parameter word bit values implemented by the Select Function (001).

The DD-40 DSU is a high performance, high capacity, fixed media magnetic storage device. Each cabinet contains four 14-inch disk drives providing a total data capacity of 5.2 gigabytes. The four individual disk drives operate together as a single unit. Each disk drive has six disk platters and ten recording surfaces. Data is accessed by 19 heads that are positioned by a precision servo mechanism. Each disk has 1420 cylinders. With the heads positioned on cylinder, 19 tracks of data are available, one track under each head. Each track contains 12 sectors. There are 512 64-bit words plus error correction data in a sector.

8.8.3 **CRAY-3 Circuits for the DD-49 and DS-40**

The CRAY-3 employs circuitry on the M module to interface between a DD-49 or DS-40 disk subsystem, two 512-word buffers and the foreground channel loop. The Foreground Processor communicates with this circuitry using five control and 16 data lines. The information is relayed on to the next node in the channel loop. A function is processed if the ID in the function word matches this controller's ID.


The DD-49/DS-40 controller communicates with the disk drive subsystem using two 24-pair cables, each with a 16-bit data path. Control signals to the drive include a 12.5 MHz Write Clock, Function Ready, and a four-bit Drive Function Code. Control signals from the drive include a 12.5 MHz Read Clock, Status/Data Ready, Error, Done, and Drive Ready. The Index/Sector Mark signal is ignored. Data is buffered between the disk subsystem and the channel loop using two alternating read/write buffers. Timing constraints within the controller and the disk drive will limit proper operation of the disk controller to a system clock in the range of 475 MHz to 525 MHz.

The controller communicates with the foreground channel loop, a single buffer stack, and the disk drive interface. It contains a controller status register, a call response register, a bus-in and drive status register, and a bus-out register. All internal data paths are 16 bits wide. The controller itself does not interpret or manipulate any disk data—it merely provides for the transfer of data between the drive and the foreground channel.

The channel interface circuitry contains the call response register and relay paths for channel loop data and control. Arriving channel loop data is made available for function parameters, drive function re-coded, drive parameters and buffer write data. Buffer readout data, status register contents and blank response data are merged for transmission to the channel loop.


The data buffers are located in a separate board stack. Each buffer has a capacity of 2048 parcels. A channel buffer pointer and a DSU buffer pointer indicate which buffer is to receive and supply data to the channel or DSU. Independent pointers and data paths enable a data transfer to or from one buffer while data is concurrently transferred from or to the other buffer. This ping-ponging of the buffers allows continuous disk read and write sequences that can avoid costly missed disk revolutions.

FIGURE 166. Block Diagram of the DD-49/DS-40 Disk Controller.



Many of the functions from the Foreground Processor contain parameters used in the disk functions. A continue-read or -write function may be pending requiring these parameters be held until the function is sent to the drive. Function parameters and write data are merged into the bus-out register for transmission over Cable A. The 8-bit controller functions are re-coded to 4-bit drive function codes and held for transmission. The module clock is divided by 40 to produce a nominal 12.5 MHz Write Clock signal to the drive. A two nanosecond pulse is provided at the active edge of Write Clock to initiate a buffer read sequence for each parcel sent during data transfers to the disk drive.

FIGURE 167. Drive Function Timing Sequence.




Notes:

1. Error and Status/Data Ready are asserted only under certain conditions.
2. Signal levels shown are for Cray-3 (GaAs) logic.

The Read Clock signal from the disk drive is re-synchronized to the 500 MHz CRAY-3 system clock by a 15-stage shift register that samples and shifts every 8ns. A two nanosecond pulse is created within the window of valid data/control signals from the drive. The Status/Data Ready, Error, and Done signals are sampled with this pulse. The bus-in and drive status registers are entered with Cable B data under certain combinations of these drive signals. Captured bus-in data is made available to each buffer stack for read function data transfers.

FIGURE 168. Cable A Timing Diagram.




Notes:

1. Write Clock is 50% duty cycle. Controller provides 80 nsec period (at 500MHz).
2. Read Clock period is 50% duty cycle.
3. Maximum 4.5 nsec rise and fall times.
4. Setup time $T_s = 20$ nsec.
5. Hold time $T_h = 20$ nsec.
6. Signals shown are Cray logic levels as measured on module.
7. Signals are inverted on cable.
8. Max potential signal skew through cabling and translators is approximately 12nsec.

FIGURE 169. Cable B Timing.

Cable B Timing (Signals from Drive)



Notes:

1. Write Clock is 50% duty cycle. Controller provides 80 nsec period (at 500MHz).
2. Read Clock period is 50% duty cycle.
3. Maximum 4.5 nsec rise and fall times.
4. Setup time $T_s = 20$ nsec.
5. Hold time $T_h = 20$ nsec.
6. Signals shown are Cray logic levels as measured on module.
7. Signals are inverted on cable.
8. Max potential signal skew through cabling and translators is approximately 12nsec.

Data transfers between the controller and the foreground channel will always be eight blocks of 256 parcels for a total of 2048 parcels. The receiving node in the transfer must be initialized first followed by the transmitting node. The receiver will immediately issue a function response of zeros. The transmitting node will then begin the transfer by sending a data pulse followed by 256 parcels of data. The receiver acknowledges receipt of the data by issuing a data pulse, after which the transmitter may begin with the next 256 parcel block. The transmitter will issue a function response of zeros at the completion of the data transfer. The primary and secondary channel transfer counters will be used to count the groups of 256 and the 256 parcels within each group, respectively.

A call pulse received during a data transfer will cause the transfer to abort and all subsequent processing will be skipped.

Data transfers between the controller and the disk drive will be either 16 or 2048 parcels long. All transfers will be in groups of 16 parcels, with an acknowledgment from the receiving device between each group of 16. Transfers to the drive will be acknowledged by the drive with the Status/Data Ready signal. Transfers from the drive will be acknowledged by the controller by issuing a Data function (0110). The primary and secondary DSU transfer counters will be used to count the groups of 16 and the 16 parcels within each group, respectively.

8.8.4 RAID Disk Array Subsystem

The RAID (Redundant Array of Inexpensive Disks) is a disk subsystem designed to provide very high data transfer rates, large capacity and reliable data availability. The subsystem is connected to the CRAY-3 computer over a full duplex HIPPI channel and is capable of sustained transfer rates of over 50 megabytes per second. The RD subsystem controller provides one or two full duplex HIPPI channel connections.

The RAID subsystem consists of a controller and one to eight racks of Disk Drawers. Each Disk Drawer contains two 1.5 gigabyte (unformatted) Head and Disk Assemblies (HDAs). A rack is made up of either 6 or 10 Disk Drawers. All racks attached to a controller must have the same number of Disk Drawers.

In the 6 Disk Drawer rack, four of the Disk Drawers contain data and the associated Error Correction Code (ECC), one Disk Drawer contains parity and one Disk Drawer is a standby. This rack provides a storage capacity of 11 gigabytes (formatted) and can sustain a transfer rate of over 25 megabytes per second.

In the 10 Disk Drawer rack, eight of the Disk Drawers contain data and the associated ECC, one Disk Drawer contains parity and one Disk Drawer is a standby. This rack provides a storage capacity of 23 gigabytes (formatted) and can sustain a transfer rate of over 50 megabytes per second.

The ECC and parity information interact to assure the integrity and availability of the data even if one disk drive in a rack fails. A failed drive can be repaired while the RAID subsystem continues to operate. The data for the failed drive is reconstructed by the subsystem as a background operation when the drive is returned to service. The individual HDAs in the subsystem are very reliable, with a MTBF of over 100,000 hours. A redundant array of these drives offers very secure storage as well as high performance.

TABLE 21. Available Models of the RD Disk Systems.

Model	No. of HIPPI Channels	No. of Disk Drawers	Transfer Rate Megabytes / Second	Capacity Gigabytes
RD106	1	6	25	11
RD206	2	6	25	11
RD110	1	10	50	23
RD210	2	10	50	23
Upgrade Racks: RD006 RD010		6 10		11 23

The capacities and transfer rates assume the use of 64K byte blocks for formatting purposes.

8.8.5 CRAY-3 HIPPI Circuitry

The CRAY-3 contains a 32-bit HIPPI source controller and a 32-bit HIPPI destination controller. Each of these controllers interfaces between a foreground channel loop and an external HIPPI simplex channel. The Foreground Processor and other I/O nodes communicate with the controllers using five control and 16 data lines each. This information is relayed to the next channel loop node. A function is processed if function group and node address fields match the controllers' group and ID.

Device control signals include CLOCK, REQUEST, CONNECT, READY, PACKET, and BURST. The HIPPI channel uses a 40ns clock for a burst bandwidth of 100 MBytes/second. Data is buffered using a 4096 x 64-bit word buffer stack and a 32 parcel FIFO. The FIFO allows sharing of the buffer stack during concurrent channel loop and HIPPI channel data transfers. One external data transfer function may be queued in the source controller to reduce idle HIPPI channel intervals due to foreground interrupt response time.

A

Integrated Circuit Packages Used in the CRAY-3

This appendix provides a listing of all the integrated circuit packages used in the CRAY-3. The table gives the name of the integrated circuit package using the latest revision available at the time of printing. It also lists every board stack where each package is used. The quantity listed is the number of times each package is used in a 16-processor CRAY-3 with the exception of those quantities followed by a single or double asterisk. The quantity followed by a single asterisk is the quantity in an M module only. The quantity followed by a double asterisk is the quantity used in a 16-processor machine but not including the packages used in any I/O modules. The last column in the table gives a very brief description of how the package is used.

TABLE 22. Integrated Circuits Used in the CRAY-3.

Package Name	Stacks Where Used	# Used (16-P)	How Used
AA-02.A	AA	512	Address Register.
AB-03.A	AA	128	Address Adder.
AC-02.A	AA	32	Adder Carry Pyramid.
AD-03.A	AB	256	Merge Data for A Registers.
AE-02.A	AB	256	Distribute Aj Data.
AF-02.A	AB	128	Distribute Ak Data.
AG-02.A	AA	16	Zero Tests.

Package Name	Stacks Where Used	# Used (16-P)	How Used
AH-03.A	AA	32	Readout Pointer.
AI-03.A	AA	32	Sequence Control.
AJ-02.A	AE-BE-CE-DE	512	Vector Element Pointer.
AK-03.A	AA	16	Sequence Control.
AL-02.A	AE-BE-CE-DE	256	Transit Time Delay.
AM-02.A	AE-BE-CE-DE	64	Common Memory Pointer.
AN-02.A	AE-BE-CE-DE	64	Vector Destination Pointer.
AO-02.A	AE-BE-CE-DE	128	Vector Length Distribution.
AP-02.A	AE-BE-CE-DE	192	Vector Unit Window.
AQ-03.A	AF-BF-CF-DF	128	First Level g Translations.
AR-03.A	AF-BF-CF-DF	64	First Level h Translations.
AS-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	1024	Vector Address Relay.
AT-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	128	Go Write Vector Register.
AU-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	128	Scalar Readout Pointers.
AV-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	256	Scalar Write Source and Destination.
AW-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	128	Vector Read Source Pointers.
AX-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	256	Vector Read Destination Pointers.
AY-02.A	AI-AJ-BI-BJ-CI-CJ-DI-DJ	128	Vector Write Source and Destination.
AZ-02.A	AM-AN-BM-BN-CM-CN-DM-DN	128	Vector Write Source and Destination.
BA-03.A	AF-BF-CF-DF	64	Go Issue Distribution.
BB-02.A	AF-BF-CF-DF	128	Go Vector Readout.
BC-02.A	AF-BF-CF-DF	64	Go Vector Write.
BD-02.A	AF	16	A Register Destination Delay.
BE-02.A	AF	16	A Register Destination Delay.
BF-03.A	AF	16	A Register Destination Delay.
BG-02.A	AF	16	Constant Data to A and S Registers.
BH-02.A	AF	16	Constant Data to A and S Registers.
BI-02.A	AF	32	Constant Data to A and S Registers.
BJ-03.A	AF	16	Address Arithmetic Control.
BK-04.A	AF-BF-CF-DF	64	Vector Function Control.
BL-02.A	AF-BF-CF	48	Functional Unit Modes.

Package Name	Stacks Where Used	# Used (16-P)	How Used
BM-02.A	AF	16	Common Memory Control.
BN-02.A	AF-BF-CF-DF	64	Local Memory Control.
BO-02.A	AF-BF-CF-DF	64	Scalar Register Readout.
BP-02.A	AF-BF-CF-DF	128	Vector Readout Source.
BQ-02.A	AF-BF-CF-DF	128	Vector Readout Destination.
BR-02.A	AF-BF-CF-DF	128	Vector Write Source and Destination.
BS-02.A	AF-BF-CF-DF	128	Scalar Write Source and Destination.
BT-02.A	AF-BF-CF-DF	64	S Register Destination Delay.
BU-02.A	AF-BF-CF-DF	64	S Register Destination Delay.
BV-02.A	AF-BF-CF-DF	64	S Register Destination Delay.
BW-02.A	AF-BF-CF-DF	64	S Register Destination Delay.
BX-02.A	AF-BF-CF-DF	64	S Register Destination Delay.
BY-03.A	AF-BF-CF-DF	64	S Register Destination Delay.
BZ-03.A	AF-BF-CF-DF	256	Constant Data Distribution.
CA-03.A	DA	16	A Register Destination Pointer.
CB-02.A	DA	16	A Register Destination Pointer.
CC-02.A	DA	16	A Register Destination Pointer.
CD-03.A	DA	16	A Register Path Conflicts.
CE-04.A	DA	16	S Register Destination Pointer.
CF-02.A	DA	16	S Register Destination Pointer.
CG-02.A	DA	16	Release S Register Reservation.
CH-03.A	DA	16	S Register Path Conflicts.
CI-04.A	DA	16	Vector Add, Shift, Logical Conflicts.
CJ-03.A	DA	16	Floating Add, Multiply Conflicts.
CK-03.A	DA	16	Branch and Local Memory Conflicts.
CL-02.A	DA	16	Common Memory Conflicts.
CM-03.A	DA	128	Instruction Parcel Selection.
CN-03.A	DA	48	Instruction ijk Distribution.
CO-02.A	DA	32	Operand Register Busy Translations.
CP-02.A	DA	128	Operand Register Conflicts.
CQ-02.A	DA	16	Instruction Group Translations.
CR-03.A	DA	16	Instruction Group Translations.

Package Name	Stacks Where Used	# Used (16-P)	How Used
CS-03.A	DA	16	Instruction Group Translations.
CT-02.A	DA	16	Instruction Group Translations.
CW-02.A	DA	64	Advance Rank Control.
CX-04.A	DA	16	Multi-Parcel Conflicts.
CY-06.A	DA	64	Valid Issue Parcel.
CZ-03.A	DA	64	Go Issue.
DA-02.A	AH and AQ	256	Front End Address Computation.
DB-02.A	AH and AQ	32	Back End Address Bits 00-03.
DB-03.A	AH, AQ and AR Revision 10 and Current Revision Stacks.	32	Back End Address Bits 00-03.
DC-02.A	AH, AQ and AR	32	Back End Address Bits 04-07.
DC-03.A	AH, AQ, and AR Revision 10 and Current Revision Stacks.	32	Back End Address Bits 04-07.
DD-02.A	AH, AQ and AR Revision 10 and Current Revision Stacks.	32	Back End Address Bits 08-11.
DD-03.A	AH, AQ and AR	32	Back End Address Bits 08-11.
DE-02.A	AH and AQ	160	Back End Address Upper Bits.
DF-02.A	AH and AQ	256	Limit Test Circuits.
DG-02.A	AH and AQ	32	Port Control.
DH-02.A	AH and AQ	32	Range Error Resolution.
DI-07.A	AL	128	Buffer Registers for Bank Select.
DJ-09.A	AL	128	Buffer Registers for Bank Select.
DK-02.A	AL	384	Buffer Registers for Bank Address.
DL-02.A	AL	32	Merge Octant Selection Data.
DM-02.A	AL	64	Merge Bank Selection Data.
DN-02.A	AL	128	Merge Bank Address Data.
DO-03.A	AL	128	Memory Octant Buffer Control.
DP-03.A	AL	16	Memory Port Control.
DQ-02.A	AP and BC	128	Buffer Element Pointer.
DR-02.A	AK and AP	128	Buffer Write Pulses.
DS-02.A	AP Replaced by EV-00.A	32	Buffer Lower Readout Data.

Package Name	Stacks Where Used	# Used (16-P)	How Used
DT-02.A	AP	112	Buffer Upper Readout Data.
DU-02.A	AP	32	Buffer Control.
DV-02.A	AK	16	Common Memory Port Control.
DW-02.A	AK	16	Memory Backup Retry Control.
DX-03.A	AK	16	Destination Tag Selection.
DY-04.A	AK	16	Scalar Path Reservation.
DZ-02.A	AK	16	Readout Buffer Entry Pointers.
EA-03.A	AK	16	Arriving Signal Delays.
EB-02.A	AK	16	Destination Tag Generator.
EC-02.A	AK	32	Merge Destination Tag and Backup.
ED-03.A	AK	256	Octant Rank Buffers.
EE-02.A	AK	64	Merge Octant Group Data.
EF-03.A	AK	64	Deliver Scalar Destination Tags.
EG-02.A	AK	64	Deliver Vector Destination Tags.
EH-02.A	AK	16	Deliver Special Destination Tags.
EI-07.A	FH-FL-GH-GL-HH-HL Replaced by ET-00.A	9216	Bank Address Registers.
EJ-07.A	FD-FP-GD-GP-HD-HP	3072	Bank Write Data Distribution.
EK-04.A	FD-FP-GD-GP-HD-HP	1536	Bank Readout Data Merge.
EM-06.A	FD-FP-GD-GP-HD-HP	512	Bank Control Sequence.
EN-03.A	BD	256	Merge Common Memory Write Data.
EO-03.A	BD	128	Form Initial Error Code.
EP-02.A	BD	64	Form Final Error Code.
EQ-06.A	FD-FP-GD-GP-HD-HP	1024	Bank Write Enable Fanout.
ER-07.A	FL-GL-HL	512	Chip Select Registers.
ES-04.A	FL-GL-HL	256	Mode Flag Registers.
ET-00.A	FH-FL-GH-GL-HH-HL Revision 10 Stacks and Current Revision Stacks. Replaces EI-07.A.	9216	Bank Address Registers for 4 Meg by 1 Memory.
EU-00.A	FD-FP-GD-GP-HD-HP Revision 10 Stacks and Current Revision Stacks. Replaces TE-05.A.	1024	Go Bank Distribution for 4 Meg by 1 Memory.

Package Name	Stacks Where Used	# Used (16-P)	How Used
EV-00.A	AP Revision 10 Stacks and Current Revision Stacks. Replaces DS-02.A.	32	Buffer Lower Readout Data for 4 Meg by 1 Memory.
FA-04.A	CA and CB	128	Write Data.
FB-03.A	CA and CB	256	Address Data.
FC-03.A	CA and CB	32	Control and Parity.
FD-03.A	CA and CB	160	Readout Data Merge.
FE-03.A	CC	64	Approximation Write Data and j Data.
FF-05.A	CC	16	Approximation Control.
FG-05.A	CC	256	Multiplicand k Data Fanout.
FH-03.A	CC	192	Multiplier j Data Fanout.
FI-03.A	CC	16	Approximation Round Summand.
FJ-03.A	CC	32	Approximation Readout Summand.
FK-03.A	CC	16	Hold Sj and Instruction Relay.
FL-03.A	CD-CH-CK	208	Multiplier First Translation.
FM-04.A	CD-CH-CK	1296	Multiplicand Re-code and Level 1 Adders.
FN-03.A	CL	16	Force Zeros.
FO-04.A	CD-CH-CK	1008	Adder Levels 2, 3, and 4.
FP-03.A	CL	64	Sign Extension.
FQ-03.A	CL	544	1-7-6 to 1-2-1-0-0 Adder.
FR-03.A	CP	64	4-4-4-4-4 to 1-2-2-2-2-1 Adder.
FS-03.A	CK and CP	32	Lower Section Carry.
FT-03.A	CP	16	Low Section Carry.
FU-03.A	CP	80	Upper Section Carry.
FV-03.A	CP	64	Front 4-Bit Add.
FW-03.A	CP	16	Round Group Final Add.
FX-02.A	CP	176	Final 4-Bit Add.
FY-02.A	CP	16	Upper Coefficient and Control.
FZ-03.A	CO and CP	192	Two-Phase Delay. Result Delay (CP).
GA-02.A	AI-AJ-AM-AN-BI-BJ-BM- BN-CI-CJ-CM-CN-DI-DJ- DM-DN	256	Scalar Add Data Group.

Package Name	Stacks Where Used	# Used (16-P)	How Used
GB-03.A	AI	16	Scalar Add Carry Propagation.
GC-03.A	BI and CI	32	Scalar Add Carry Propagation.
GD-03.A	DI	16	Scalar Add Carry Propagation.
GE-02.A	AO	128	Front Portion Scalar Shift.
GF-02.A	AO	128	Back Portion Scalar Shift.
GG-02.A	AO	128	Merge Scalar Shift Results.
GH-03.A	AO	16	Scalar Shift Count.
GI-03.A	AO	16	Scalar Shift Control.
GJ-02.A	AO	64	Front Section Scalar Population Count.
GK-02.A	AO	64	Front Section Scalar Leading Zero Count.
GL-02.A	AO	16	Center Section Scalar Population Count.
GM-02.A	AO	16	Back Section Scalar Leading Zero Count.
GN-02.A	AO	16	Merge Population and Leading Zero Count.
GO-02.A	AF and DF	64	Parcel Delay Circuits.
GP-02.A	DF	16	Foreground Interface Signals.
GQ-02.A	DF	32	Constant Data to S Registers.
GR-02.A	DF	16	Functional Unit Modes.
GS-02.A	DF	16	Go Vector Unit Windows.
GT-02.A	AH and AQ	16	Delay Registers for Constant Data.
GU-03.A	AB	16	Selection Control.
GW-02.A	CF	64	Constant Data and Aj Sign Distribution.
GX-02.A	CF	16	Floating Multiply Control.
GY-02.A	BF	16	Floating Add and Common Memory Control.
GZ-03.A	BF	16	Branch Control.
HA-03.A	DH	16	Channel Control Signals.
HB-03.A	DH	64	Channel Data Node.
HC-02.A	DH	16	Early Channel Functions.
HD-03.A	DH	16	Real Time Clock and Instruction 037.

Package Name	Stacks Where Used	# Used (16-P)	How Used
HE-03.A	DH	16	Group A Registers Bits 00-07.
HF-03.A	DH	16	Group A Registers Bits 08-15.
HG-02.A	DH	16	Group A Registers Bits 16-23.
HH-02.A	DH	16	Group A Registers Bits 24-31.
HI-03.A	DH	48	Group B Registers General.
HJ-03.A	DH	16	Group B Registers Bits 08-15
HK-02.A	DH	128	Data Distribution.
HL-04.A	DH	16	Semaphore Flags.
HM-02.A	DH	16	Channel Functions for Registers.
HN-03.A	DH	16	Channel Functions and Memory Slot.
HP-02.A	DD	64	Buffer Address Distribution.
HQ-03.A	DD	128	Buffer Write Pulses.
HR-02.A	DD	128	Buffer Readout Distribution.
HS-02.A	DD	16	Buffer Address Pointer.
HT-02.A	DD	16	Buffer Length Counter.
HU-02.A	DD	64	Memory Address Register.
HV-02.A	DD	16	Reference Length Register.
HW-02.A	DD	16	Channel Function Translations.
HX-02.A	DD	16	Buffer Control Sequences.
HY-02.A	DD	16	Destination Selection Control.
IA-02.A	BA	128	Lower Six Bits Coincidence Test.
IB-02.A	BA	384	Eight Bits Coincidence Test.
IC-02.A	BA	32	Merge Coincidence Tests.
ID-03.A	BB	128	Branch Address Register.
IE-03.A	BB	16	Instruction Word Counter.
IF-03.A	BB	112	Instruction Word Counter.
IG-02.A	BB	16	Word Counter Carry Circuits.
IH-03.A	BB	16	Program Address Register.
II-03.A	BB	80	Program Address Register.
IJ-03.A	BB	128	Memory Address Merge.
IK-02.A	BB	16	Current Field Pointer.
IL-03.A	BB	16	Instruction Stack Readout Pointer.

Package Name	Stacks Where Used	# Used (16-P)	How Used
IM-02.A	BB	16	Test Branch Address.
IN-02.A	BB	16	A Register Branches.
IO-02.A	BB	16	Semaphore Branches.
IP-03.A	BB	16	S Register Branches.
IQ-02.A	BB	16	Program Address Parcel Pointer.
IR-02.A	BB	16	Foreground Sequences.
IS-03.A	BB	16	Advance Buffer Field.
IT-03.A	BB	16	Common Memory Requests.
JA-02.A	MC	1*	Foreground Channel Control.
JB-01.A	MC	4*	Foreground Channel Data.
JC-01.A	MC	4*	Buffer Address and Write Data.
JD-01.A	MC	4*	Channel Data Merge.
JE-01.A	MC	1*	Function Translation.
JF-01.A	MC	1*	Channel Transfer Length.
JG-01.A	MC	1*	Buffer Stack Control.
JI-01.A	MC	3*	Ready Counters and Adder.
JJ-01.A	MC	1*	Data FIFO Address.
JK-01.A	MC	8*	DATA BUS.
JL-01.A	MC	2*	PARITY BUS and DATA Errors.
JM-02.A	MC	1*	CONNECT and READY.
JN-02.A	MC	1*	REQUEST, PACKET, and BURST.
JO-01.A	MC	2*	Packet File Data and Address.
JP-01.A	MC	1*	End Packet Ranks.
JQ-01.A	MC	1*	Packet File Count and Control.
JZ-01.A	MC	2*	CLOCK Edge Detect and Fanout.
KA-01.A	MD	1*	Foreground Channel Control.
KB-01.A	MD	4*	Foreground Channel Data.
KC-01.A	MD	4*	Buffer Address and Write Data.
KD-01.A	MD	4*	Buffer Readout Data.
KE-01.A	MD	1*	Function Translation.
KF-01.A	MD	1*	Channel Transfer Length.
KG-01.A	MD	1*	Buffer Stack Control.

Package Name	Stacks Where Used	# Used (16-P)	How Used
KH-01.A	MD	1*	Device Transfer Length.
KI-01.A	MD	1*	Packet Burst and Ready Count.
KJ-01.A	MD	1*	FIFO Address and Count.
KK-01.A	MD	8*	DATA BUS.
KL-01.A	MD	1*	PARITY BUS.
KM-01.A	MD	1*	CONNECT and READY.
KN-01.A	MD	1*	REQUEST, PACKET, and BURST.
KO-01.A	MD	1*	BURST Limit.
KP-01.A	MD	1*	PACKET and BURST Control.
KZ-01.A	MD	2*	CLOCK Edge Detect and Fanout.
LA-07.A	EC	4	Address and Control.
LB-05.A	EC	4	Write Data.
LC-03.A	EC	4	Readout Data.
LD-02.A	EG	8	Register Adder.
LE-02.A	EG	2	Adder Carry Pyramid.
LF-02.A	EG	4	Functional Unit Merge.
LG-02.A	EG	4	Real Time Clock.
LH-02.A	EG	4	Front End Shift.
LI-02.A	EG	2	Back End Shift.
LJ-02.A	EH	4	A and B Register Function Data.
LK-04.A	EH	16	Channel Response Register.
LL-02.A	EH	4	Channel Read Merge.
LM-03.A	EH	2	Channel Control.
LN-06.A	EK	1	Re-sync Counter.
LO-07.A	EK	1	Console Control.
LP-03.A	EL	2	Clear Real Time Clock and System Deadstart Fanout. Force One Fanout.
MA-02.A	EK	3	Instruction Address Register.
MB-02.A	EK	2	Branch Address Merge.
MC-06.A	EF and EJ	16	Bank Address and Control.
MD-05.A	EF and EJ	2	Bank Write Data.
ME-03.A	EF and EJ	8	Readout Data.

Package Name	Stacks Where Used	# Used (16-P)	How Used
MF-04.A	EF and EJ	2	Bank and Group Selection.
MG-04.A	EK	4	Console Assembly Register.
MH-07.A	EK	4	A, B Register Disassembly Register.
MI-03.A	EK	4	Constant Data Assembly.
MJ-02.A	EK	1	A Register Instruction Group Translations.
MK-03.A	EK	1	B Register Instruction Group Translations.
ML-03.A	EK	1	Instruction Group Translations (Branch/Multi-Byte).
MM-02.A	EK	1	Go Issue and Rank One Valid Translation.
MN-04.A	EK	1	A, B Register and Local Memory Conflicts.
MO-02.A	EK	1	A, B Register Path Conflicts.
MP-03.A	EK	1	Console Control and Branch Conflicts.
MQ-02.A	EK	1	Channel Control.
MR-02.A	EG	16	A, B, and C Register.
MS-02.A	EG	8	Register Data Merge.
MT-02.A	EK	1	B Register Branch Control.
MU-02.A	EK	1	Functional Unit Control.
MV-02.A	EK	1	Constant Control.
MW-03.A	EK	1	A, B Register Destination Chain.
MX-02.A	EK	1	A, B Register Destination Chain.
MY-02.A	EK	1	A, B Register Destination Pointer.
MZ-02.A	EK	1	Control Relay.
NA-03.A	MK	1*	Foreground Channel Control.
NB-02.A	MK	4*	Foreground Channel Data.
NC-02.A	MK	1*	Function Translations.
ND-03.A	MK	2*	Merge Data to Foreground.
NE-02.A	MK	1*	Dump Input Buffer.
NF-02.A	MK	1*	Load Input Buffer.
NG-02.A	MK	1*	Function Translations.

Package Name	Stacks Where Used	# Used (16-P)	How Used
NH-02.A	MK	2*	Buffer Length Counter.
NI-04.A	MK	1*	Load Output Buffer.
NJ-03.A	MK	1*	Device Signal Generator.
NK-03.A	MK	1*	Dump Output Buffer.
NL-02.A	MK	4*	Arriving Device Data.
NM-02.A	MK	1*	Output Data to Device.
NN-03.A	MK	2*	Synchronize Inputs.
NO-00.A	MG and MO	8*	Bank Address.
NP-01.A	MG and MO	2*	Sequence Control.
NQ-01.A	MG and MO	2*	Write Data and Parity.
NR-01.A	MG and MO	2*	Readout Data.
NS-03.A	ED	1	Channel One Control.
NT-03.A	ED	1	Channel Zero Control.
NU-03.A	ED	1	Readout Delay Chain.
NV-02.A	ED	4	Channel One Input Data.
NW-04.A	ED	4	Channel One Output Data.
NX-03.A	ED	8	Channel Zero, Two Input Data.
NY-02.A	ED	2	Merge Data to Memory.
NZ-03.A	ED	4	Channel Zero Output Data.
PA-02.A	DB	96	Buffer Address Distribution.
PB-03.A	DB	256	Buffer Write Pulses.
PC-02.A	DB	128	Buffer Readout Distribution.
PD-02.A	DC	16	Destination Tag to Stack Buffers.
PE-02.A	DC	16	Destination Tag to Parity Buffers.
PF-02.A	DC	16	Parity Buffer Write Pulses.
PG-02.A	DC	128	Valid Word Register.
PH-03.A	DC	16	Stack Read Address Register.
PI-03.A	DC	16	Instruction Stack Sequence Control.
PJ-02.A	DC	16	Form Stack Parity Error Signals.
PP-02.A	AD	48	Two Bits Vector Length Distribution.
PQ-02.A	AC	256	Pyramid Formation.
PR-03.A	AC	32	Pyramid Formation.

Package Name	Stacks Where Used	# Used (16-P)	How Used
PS-02.A	AC	128	Operand Distribution.
PT-02.A	AD	64	Three Input Adders 2-3-2-3-2-3.
PU-02.A	AD	16	Three Input Adders 2-3-3-4-3-4.
PV-02.A	AD	320	Three Input Adders 4-4-4-5.
PW-02.A	AD	32	Carry Propagation.
PX-03.A	AD	128	Four Bit Slice Final Add.
PY-03.A	AD	16	Vector Length Register.
QA-03.A	FL	128	Merge Bank Pointer.
QB-03.A	FH and FL	640	Group Bank Translation. Pointer Translation (FL).
QC-03.A	FH	512	Bank Reservation Flags.
QD-03.A	FH	96	Test for Duplicate Request.
QE-03.A	FL	128	Test Bank Busy.
QF-03.A	FL	32	Processor Resumes.
QG-03.A	FL	256	Go Bank Generation.
QH-03.A	FL	96	Distribute Bank Pointer.
QI-03.A	FL	64	Distribute Pointer Valid.
QK-04.A	GH and GL	1280	Merge Write Data Packet.
QL-05.A	GD and GP	640	Group Element Fanout.
QM-03.A	GD and GP	2560	Merge Write Data for Four Banks.
QN-04.A	GD and GP	2560	Bank Pointer Translation.
QP-03.A	GH and GL	320	Group Secondary Rank Valid and Retry.
QQ-04.A	HD and HP	2304	Merge Bank Readout Data.
QR-03.A	HL	576	Final Readout Merge.
QS-04.A	FL	64	Bank Error Pointer.
QT-03.A	BD	16	Memory Write Merge Control.
QU-02.A	BC	16	Retry Buffer Control.
QV-02.A	BH and BL	128	Write Buffer Control.
QW-03.A	HH	576	Group Bank Pointer.
QX-03.A	HH	144	Control Signals.
RA-02.A	BP	64	Front End of Exponent Difference.
RB-02.A	BP	64	Back End of Exponent Difference.

Package Name	Stacks Where Used	# Used (16-P)	How Used
RC-02.A	BP	64	Final Exponent Selection.
RE-02.A	BP	64	Front End of Exponent Adjustment.
RF-02.A	BP	16	Coefficient Shift Selection.
RG-02.A	BP	48	Coefficient Shift Bit Relay.
RH-02.A	BK	16	Miscellaneous Control.
RI-03.A	BO	16	Instruction Entry and Fanout.
RJ-02.A	BO and BP	48	Instruction Decode.
RK-02.A	BK	16	Leading Zero Count Section Translations.
RL-02.A	BK	48	Leading Zero Count Bit Translations.
RM-02.A	BP	16	Normalization Shift Count.
RN-02.A	BP	16	Exponent Control.
RP-02.A	BP	16	Exponent Control, Sign Bit, and Range Error.
RQ-02.A	BO	192	Scalar/Vector Entry.
RS-02.A	BO and BP	256	Front Coefficient Shift. Front End of Coefficient Normalization (BP).
RT-02.A	BO	64	Unshifted Data Delay.
RU-02.A	BK	192	Front Section of Coefficient Adder.
RV-02.A	BK	96	Carry Propagation.
RW-02.A	BK	192	Back Section of Coefficient Adder.
RX-02.A	BK	96	Shift 0, -1
RY-02.A	BP	96	Back End of Coefficient Normalization.
RZ-02.A	BP	64	Back End of Exponent Adjustment.
SC-00.A	AG-BG-CA-CB-CG-DG- EC-EF-EJ-MA-MB-MG- MM-MN-MO	7348**	Local Memory.
SD-00.A	SA-SE-SM	147456	Common Memory.
SF-00.A	SB-SF-SN	73728	Common Memory.
SM-00.A	XA-XE-XM	36864	Common Memory.
TC-03.A	HH	432	Process Group Address Data.
TD-03.A	HD and HP	1728	Bank Address Distribution.

Package Name	Stacks Where Used	# Used (16-P)	How Used
TE-05.A	FD-FP-GD-GP-HD-HP Replaced by EU-00.A	1024	Go Bank Distribution.
TF-03.A	FD	32	Slot Time Distribution.
TG-03.A	FD and FP	128	Bank Pointer Distribution.
TH-02.A	BC	288	Merge Write Data.
TI-04.A	BH and BL	1152	Octant Data Buffers.
TJ-05.A	GH and GL	960	Group Bank Destination Pointer.
TK-03.A	FD	64	Group Parameter Distribution.
TL-01.A	MH-ML-MP	3*	Channel Control.
TM-01.A	MH-ML-MP	12*	Channel Data.
TN-01.A	MH-ML-MP	3*	Function Translation.
TO-01.A	MH-ML-MP	3*	Function Re-code.
TP-00.A	MH-ML-MP	3*	Channel to Buffer Control.
TQ-00.A	MH-ML-MP	3*	DSU to Buffer Control.
TR-00.A	MH-ML-MP	12*	Buffer Merge.
TS-00.A	MH-ML-MP	24*	FIFO.
TT-00.A	MH-ML-MP	3*	Write Clock.
TU-00.A	MH-ML-MP	3*	DSU Read Control.
TV-00.A	MH-ML-MP	3*	DSU Write Control.
TW-00.A	MH-ML-MP	12*	DSU Bus-out Data.
TX-00.A	MH-ML-MP	3*	Read Clock Re-synchronization.
TY-00.A	MH-ML-MP	3*	DSU Input Control.
TZ-00.A	MH-ML-MP	6*	DSU Bus-in Data.
UA-02.A	DO	64	Front Section Population Count.
UB-02.A	DO	16	Center Section Population Count.
UC-02.A	DO	16	Lowest Order Four Bits of Result.
UD-02.A	DO	16	Bits 04 through 07 of Result.
UE-02.A	DO	64	Front Section of Leading Zero Count.
UF-02.A	DO	16	Back Section of Leading Zero Count.
UG-02.A	DO	128	Delay Registers for Vector Shift.
UH-02.A	DO	128	Front Section of Vector Shift.
UI-02.A	DO	128	Back Section of Vector Shift.

Package Name	Stacks Where Used	# Used (16-P)	How Used
UJ-02.A	DO	16	Vector Shift Count.
UK-02.A	DO	16	Sequence Control.
UL-02.A	DO	112	Bits 08 through 63 of Result.
UM-00.A	MH-ML-MP	12*	Status Registers.
UO-00.A	MH-ML-MP	6*	Asynchronous Latch.
VA-02.A	AN-BN-CN-DN	256	Vector Add Front Section.
VB-02.A	AN-BN-CN-DN	256	Vector Add Back Section.
VC-02.A	BN-CN-DN	48	Vector Add Carry Propagation.
VD-02.A	AN-BN-CN-DN	64	Compress Iota Mask Register.
VE-03.A	AN-BN-CN-DN	64	Vector Add Control Circuits.
VF-02.A	DN	16	Compress Iota Control.
VG-02.A	AM-BM-CM-DM	256	Vector Logical Data Merge.
VH-02.A	AM-BM-CM-DM	256	Vector Mask Data Registers.
VI-02.A	AM-BM-CM-DM	128	Real Time Clock Register.
VJ-02.A	AM-BM-CM-DM	64	Vector Logical Boundary Circuits.
VK-02.A	AM-BM-CM-DM	64	Vector Mask Control.
VL-03.A	AM-BM-CM-DM	64	Vector Logical Control.
VP-03.A	AI-AJ-AM-AN-BI-BJ-BM-BN-CI-CJ-CM-CN-DI-DJ-DM-DN	2048	Vector Address Register.
VQ-02.A	AI-AJ-AM-AN-BI-BJ-BM-BN-CI-CJ-CM-CN-DI-DJ-DM-DN	2048	Merge Data to Vector Registers.
VR-01.A or VR-01.B	AI-AJ-AK-AM-AN-AP-BC-BI-BJ-BM-BN-CI-CJ-CM-CN-DB-DC-DI-DJ-DM-DN-MC-MD	3232**	Vector Registers. Buffer Registers (AK- AP-BC-DB-DC). FIFO Buffer Data (MC-MD).
VS-02.A	AI-AJ-AM-AN-BI-BJ-BM-BN-CI-CJ-CM-CN-DI-DJ-DM-DN	1024	Scalar Registers.
VT-02.A	AI-AJ-AM-AN-BI-BJ-BM-BN-CI-CJ-CM-CN-DI-DJ-DM-DN	3072	Distribute Vector Register Data.
VU-02.A	AI-AJ-AM-AN-BI-BJ-BM-BN-CI-CJ-CM-CN-DI-DJ-DM-DN	256	Scalar Non-Zero Test.

Package Name	Stacks Where Used	# Used (16-P)	How Used
VV-01.A	DD-DK-DL-DP	592	Buffer Registers. Identical to VR-01.A except for inverted clock.
VX-02.A	AE-BE-CE-DE	512	Vector Read Length Counter.
VY-02.A	AE-BE-CE-DE	512	Vector Write Length Counter.
VZ-02.A	AN	16	Vector Add Carry Propagation.
WA-02.A	AG-BG-CG-DG	128	Address Data Merge.
WB-02.A	AG-BG-CG-DG	128	Write Data Merge.
WC-06.A	AG-BG-CG-DG-MA-MB-MM-MN	512	Bank Address and Control.
WD-05.A	AG-BG-CG-DG-MA-MB-MG-MM-MN-MO	256	Bank Write Data.
WE-02.A	AG-BG-CG-DG-MA-MB-MG-MM-MN-MO	256	Readout Data.
WF-02.A	AG-BG-CG-DG	64	Bank Selection.
WG-02.A	AG-BG-CG-DG	64	Parity Error Check.
WH-02.A	MA-MB-MM-MN	20*	Address Register.
WI-02.A	MA-MB-MM-MN	10*	Bank Control.
WJ-02.A	MA-MB-MM-MN	10*	Control and Parity Check.
WK-02.A	MA-MB-MM-MN	20*	Write Data Delay and Parity.
WL-02.A	CL-MD	32**	Center Mat Constants. Node Address Constants (MD).
XI-02.A	DL and DP	320	Merge of Packets 0-3 from Octants. Merge of Check Bits and Bank Address (DP).
XJ-03.A	DL	32	Control Relay.
XK-02.A	DK and DL	224	Front of SEC-DED (Bypass and Buffer).
XL-02.A	DL	128	Back of SEC-DED for All Packets.
XM-02.A	DL	64	Front of SEC-DED for Packet 3 (Direct).
XN-02.A	DK and DP	256	Path Selection.
XO-02.A	DL and DP	96	Merge of Group B Packet 3 from Octants. Path Selection (DP).
XP-02.A	DL and DP	160	Merge and Data Delay. Read Data Merge (DP).

Package Name	Stacks Where Used	# Used (16-P)	How Used
XQ-02.A	DK-DL-DP	240	Merge and Data Delay. Check Bit and Bank Address Merge (DP).
XR-02.A	DP	64	Syndrome Code Translation.
XS-02.A	DP	32	Syndrome Code Translation.
XT-02.A	DP	16	Error Detection Translation.
XU-02.A	DP	256	Corrected Read Data Output.
XV-02.A	DK-DL-DP	48	Control for Buffer Address and Path Selection.
XW-02.A	DK-DL-DP	96	Path Selection Control.
XX-02.A	DL	16	One Valid Reference Translation.
XY-02.A	DP	16	Octant Flag Merge.
XZ-02.A	DK-DL-DP	272	Buffer Address Register.
YA-03.A	CO	64	Exponent Operand Data.
YB-03.A	CO	64	Summation and Approximation Shift.
YC-03.A	CO	64	Final Exponent.
YD-03.A	CO	16	Instruction Decode.
YE-04.A	CO	16	Instruction Mode Delay Chain.
YF-03.A	CO	16	Operand Conditions.
YG-03.A	CO	16	Overflow/Underflow.
ZA-01.A	AA-AB-AD-AE-AF-AG-AI- AJ-AK-AL-AM-AN-AO- AP-AQ-BA-BB-BC-BD- BE-BF-BG-BH-BI-BJ-BK- BL-BM-BN-BO-BP-CA- CB-CC-CD-CE-CF-CG- CI-CJ-CK-CH-CL-CM- CN-CO-CP-DA-DB-DC- DD-DE-DF-DG-DH-DI- DJ-DK-DL-DM-DN-DO- DP-EA-EB-EC-ED-EE- EF-EG-EH-EI-EJ-EK-EL- EM-EN-EO-EP-FD-FH- FL-FP-GD-GH-GL-GP- HD-HH-HL-HP-MA-MB- MC-MD-ME-MF-MG-MH- MI-MJ-MK- ML-MN-MO-MP	2370**	Clock Distribution.

Package Name	Stacks Where Used	# Used (16-P)	How Used
ZT-01.A	AG-BG-CA-CB-CG-DG- EC-EF-EJ-EK-MA-MB- MF-MG-MH-MJ-MK-ML- MN-MO-MP	495**	Termination Resistors.

Appendix

B

CRAY-3 Board Stack Assignments

This appendix contains a series of charts which summarize the assignment of the different logical functions of the CRAY-3 among the board stacks of each module. The integrated circuit packages which are assigned to each of the four boards for each board stack are also indicated by their two letter designations. The two letter designations, which represent each individual die, are placed in their proper orientation on each board just as you would find the die themselves.

A Module Background Processor

Address Registers Address Adder		Address Data Merge Address Data Distribution		Address Multiply Pyramid Formation		Address Multiply Product Summation	
A		B		C		D	
Vector Length Register Octant Clock Distribution							
1 AB AH AA AA -- -- AA AA AC -- AA AA AB AI AA AA	2 AB AH AA AA -- -- AA AA AG -- AA AA AB AI AA AA	1 AD AE AF -- AD AE -- -- AD AE -- -- AD AE AF GU	2 AD AE AF -- AD AE -- -- AD AE -- -- AD AE AF --	1 PS -- PQ PQ PS -- PQ PR PS -- PQ PR PS -- PQ --	2 -- -- PQ PQ -- -- PQ PQ -- -- PQ PQ -- -- PQ --	1 PV PV -- -- PV PV PV PV PV PV PV PV PV PV PV --	2 -- PV PV PT -- PV PV PT -- PV PV PT -- PV PU PT
3 AB -- AA AA -- -- AA AA AC -- AA AA AB AK AA AA	4 AB -- AA AA -- ZA AA AA -- -- AA AA AB -- AA AA	3 AD AE AF -- AD AE -- -- AD AE -- -- AD AE AF --	4 AD AE AF -- AD AE ZA -- AD AE -- -- AD AE AF --	3 -- -- PQ -- -- -- PQ -- -- ZA -- -- -- PQ --	4 PS -- PQ PQ PS -- PQ PQ PS -- PQ PQ PS -- PQ --	3 -- -- -- -- PP PP PP -- -- PY ZA -- -- -- -- --	4 -- PW PX PX -- -- PX PX -- PW PX PX -- PX PX
Vector Control Vector Length Window		Instruction Parcel Translation		Local Memory 00-15		Common Memory Address Arithmetic	
E		F		G		H	
Base Address Register Limit Address Register Module Clock Distribution							
1 -- -- -- -- -- -- AO AP AL AL -- -- VY VY VY VY	2 -- -- -- -- -- -- AN AP AL AL -- -- VY VY VY VY	1 BG BH BI BI BC -- GO EM BO -- -- BK BP BQ BP BQ	2 BJ BD BE BF BB -- -- BN BB -- -- BL BR BS BR BS	1 SC SC SC SC WE WE WE WE WD WC WC SC SC SC SC SC	2 SC SC SC SC ZT WB WB ZT WD WC WC SC SC SC SC SC	1 DG -- DF DF DF DF DA DA -- ZA DA DA DB DC DC DB	2 DG -- DF DF DF DF DA DA -- ZA DA DA DD DE DE DD
3 -- -- -- -- -- -- AM -- VX VX VX VX AJ AJ AJ AJ	4 -- -- -- -- -- ZA AO -- VX VX VX VX AJ AJ AJ AJ	3 AQ AQ BZ BZ -- -- -- AR BT -- -- BY BU BV BW BX	4 -- BA BZ BZ -- -- -- -- -- ZA -- -- -- -- -- --	3 SC SC SC SC ZT WF WG ZT WD WC WC SC SC SC SC SC	4 SC SC SC SC WA WA ZA ZT WD WC WC SC SC SC SC SC	3 GT -- DF DF DF DF DA DA -- DH DA DA DE DE DE DE	4 -- -- DF DF DF DF DA DA -- DH DA DA DE DE DE DE
Scalar Registers 00-03 Scalar Logical 00-03 Scalar Add 00-03 Vector Registers 00-03		Scalar Registers 04-07 Scalar Logical 04-07 Scalar Add 04-07 Vector Registers 04-07		Common Memory Destination Tags		Common Memory Address Distribution	
I		J		K		L	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 ED EG -- EH ED EF -- -- ED EE DZ DR ED VR VR VR	2 ED EG -- -- ED EF -- EA ED EE -- EB ED -- EC EC	1 DO DO DK DK DN DL DK DK DN DM DK DK DI DI DJ DJ	2 DO DO DK DK DN DL DK DK DN DM DK DK DI DI DJ DJ
3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY -- -- ZA -- -- GA VU GB	3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY -- -- ZA -- VU GA -- --	3 ED EG -- DV ED EF -- DW ED EE -- DX ED -- -- DY	4 ED EG -- -- ED EF ZA -- ED EE -- -- ED -- -- --	3 DO DO DK DK DN DP DK DK DN DM DK DK DI DI DJ DJ	4 DO DO DK DK DN ZA DK DK DN DM DK DK DI DI DJ DJ
Scalar Registers 08-11 Scalar Logical 08-11 Scalar Add 08-11 Vector Registers 08-11 Vector Logical 00-15 Vector Mask 00-15 Real Time Clock 00-15		Scalar Registers 12-15 Scalar Logical 12-15 Scalar Add 12-15 Vector Registers 12-15 Vector Add 00-15 Comp Iota Mask 00-15		Scalar Shift Scalar Population Count Scalar Leading Zero Count		Common Memory Address Retry Buffers	
M		N		O		P	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 GF GF GH GI -- -- -- -- GK GE -- -- GJ GE GG GG	2 GF GF GM GN -- -- -- -- GK GE -- -- GJ GE GG GG	1 DQ -- -- DQ VR -- -- VR VR DS DS VR VR DR DR VR	2 DQ -- -- DQ VR DT DT VR VR DT DT VR VR DR DR VR
3 VS VS VS VS VJ VK VL AZ -- -- -- -- VT VT VT VT	4 VG VG VG VG VH VH VH VH VI -- ZA VI -- GA VU --	3 VS VS VS VS -- -- -- AZ -- -- ZA -- VT VT VT VT	4 VA VA VA VA VB VB VB VB VZ VD VE -- -- GA VU --	3 GF GF GL -- -- -- -- -- GK GE -- -- GJ GE GG GG	4 GF GF -- -- -- ZA -- -- GK GE -- -- GJ GE GG GG	3 -- -- -- -- VR -- -- VR VR DT DT VR -- DR DR VR	4 -- -- -- -- DU ZA -- DU -- DT VR -- DR VR

B Module

Background Processor

Branch Control Coincidence Test Program Register 00-07 16-23		Branch Control Program Register 08-11 24-31		Common Memory Write Write Retry Buffer		Common Memory Write Error Correction Code	
1 -- IA IA	2 -- IA IA	1 ID IH II II	2 ID II II II	1 VR VR VR --	2 VR VR VR --	1 -- EO EN --	2 -- EP EN --
-- IC IB IB	-- IC IB IB	ID -- IJ IJ	ID -- IJ IJ	VR VR VR --	VR VR VR --	-- EO EN --	-- EP EN --
-- IB IB	-- IB IB	ID IK IJ IJ	ID IL IJ IJ	VR VR VR --	VR VR VR --	-- EO EN --	-- EP EN --
-- IB IB	-- IB IB	ID -- IN --	ID -- IO --	-- DQ DQ --	-- DQ DQ --	-- EO EN --	-- EP EN --
3 -- IA IA	4 -- IA IA	3 IE -- --	4 IF -- --	3 -- TH TH TH	4 -- QU TH TH	3 -- EO EN --	4 -- -- EN --
-- IB IB	-- IB IB	IF -- -- IR	IF -- IS IT	-- TH TH	-- TH TH	-- EO EN --	-- QT EN --
-- IB IB	-- ZA IB IB	IF IG --	IF IM ZA	-- TH TH	-- ZA TH TH	-- EO EN --	-- ZA EN --
-- IB IB	-- IB IB	IF -- IP --	IF -- IQ --	-- TH TH TH	-- TH TH	-- EO EN --	-- -- EN --
Vector Control Vector Length Window		Instruction Parcel Translation		Local Memory 16-31		Common Memory Write Lower Octant Buffers	
1 -- AO AP	2 -- AN AP	1 -- GY --	2 -- GZ --	1 SC SC SC SC	2 SC SC SC SC	1 -- TI TI	2 -- TI TI
AL AL	AL AL	BC -- --	BB -- -- BN	WE WE WE WE	ZT WB WB ZT	-- QV TI TI	-- QV TI TI
VY VY VY VY	VY VY VY VY	BO -- -- BK	BB -- -- BL	WD WC WC SC	WD WC WC SC	-- TI TI	-- TI TI
		BP BQ BP BQ	BR BS BR BS	SC SC SC SC	SC SC SC SC	-- TI TI TI	-- TI TI TI
3 -- AM --	4 -- ZA AO --	3 AQ AQ BZ BZ	4 -- BA BZ BZ	3 SC SC SC SC	4 SC SC SC SC	3 -- TI TI	4 -- TI TI
VX VX VX VX	VX VX VX VX	-- -- AR	-- --	ZT WF WG ZT	WA WA ZA ZT	-- QV TI TI	-- QV TI TI
AJ AJ AJ AJ	AJ AJ AJ AJ	BT -- -- BY	-- ZA --	WD WC WC SC	WD WC WC SC	-- TI TI	-- ZA TI TI
		BU BV BW BX	-- --	SC SC SC SC	SC SC SC SC	-- TI TI TI	-- TI TI TI
Scalar Registers 16-19 Scalar Logical 16-19 Scalar Add 16-19 Vector Registers 16-19		Scalar Registers 20-23 Scalar Logical 20-23 Scalar Add 20-23 Vector Registers 20-23		Floating Add Coefficient Add Leading Zero Count Back Section Coefficient Shift		Common Memory Write Upper Octant Buffers Module Clock Distribution	
1 VQ VQ VQ VQ	2 VQ VQ VQ VQ	1 VQ VQ VQ VQ	2 VQ VQ VQ VQ	1 -- -- --	2 -- -- --	1 -- TI TI	2 -- TI TI
VR VR VR VR	VR VR VR VR	VR VR VR VR	VR VR VR VR	RW RW RW RX	RW RW RW RX	-- QV TI TI	-- QV TI TI
VP VP VP VP	VP VP VP VP	VP VP VP VP	VP VP VP VP	-- -- RX	RV RV RV RX	-- TI TI	-- TI TI
VT VT VT VT	VT VT VT VT	VT VT VT VT	VT VT VT VT	RU RU RU RL	RU RU RU RL	-- TI TI TI	-- TI TI TI
3 VS VS VS VS	4 AS AS AS AS	3 VS VS VS VS	4 AS AS AS AS	3 -- -- --	4 -- -- --	3 -- TI TI	4 -- TI TI
AT AU AV AV	AW AX AX AY	AT AU AV AV	AW AX AX AY	RW RW RW RX	RW RW RW RH	-- QV TI TI	-- QV TI TI
AS AS AS AS	-- ZA --	AS AS AS AS	-- ZA --	RV RV RV RX	-- ZA -- RK	-- ZA TI TI	-- ZA TI TI
VT VT VT VT	-- GA VU GC	VT VT VT VT	VU GA --	RU RU RU RL	RU RU RU --	-- TI TI TI	-- TI TI TI
Scalar Registers 24-27 Scalar Logical 24-27 Scalar Add 24-27 Vector Registers 24-27 Vector Logical 16-31 Vector Mask 16-31 Real Time Clock 16-31		Scalar Registers 28-31 Scalar Logical 28-31 Scalar Add 28-31 Vector Registers 28-31 Vector Add 16-31 Comp Iota Mask 16-31		Floating Add Coefficient Entry Coefficient Shift		Floating Add Exponent Processing Coefficient Normalization	
1 VQ VQ VQ VQ	2 VQ VQ VQ VQ	1 VQ VQ VQ VQ	2 VQ VQ VQ VQ	1 RS RS RS RS	2 RS RS RS RS	1 RG RC RC --	2 RG RC RC --
VR VR VR VR	VR VR VR VR	VR VR VR VR	VR VR VR VR	-- -- --	-- -- --	RB RS RS RE	RB RS RS RE
VP VP VP VP	VP VP VP VP	VP VP VP VP	VP VP VP VP	-- -- --	-- -- --	-- -- --	-- -- --
VT VT VT VT	VT VT VT VT	VT VT VT VT	VT VT VT VT	RQ RQ RQ RQ	RQ RQ RQ RQ	RA RY RY RZ	RA RY RY RZ
3 VS VS VS VS	4 VG VG VG VG	3 VS VS VS VS	4 VA VA VA VA	3 RT RT RT RT	4 -- -- --	3 -- -- --	4 RG RM RN RP
VJ VK VL AZ	VH VH VH VH	-- -- AZ	VB VB VB VB	-- -- --	-- ZA RI	RB RS RS RE	RB RS RS RE
-- -- --	VI -- ZA VI	-- ZA --	VC VD VE --	-- -- --	-- RJ RJ	-- RF ZA	-- RJ --
VT VT VT VT	-- GA VU --	VT VT VT VT	-- GA VU --	RQ -- -- RQ	RQ -- -- RQ	RA RY RY RZ	RA -- -- RZ

C Module

Background Processor

Floating Multiply

Reciprocal Approximation A		Reciprocal Approximation B		Floating Multiply Approximation Coefficient Distribution C		Floating Multiply Segment Three 47-32 D	
1 SC SC SC SC SC ZT FB FA SC ZT FB FC SC SC SC SC	2 SC SC SC SC SC ZT FB FA SC FD FB FD SC SC SC SC	1 SC SC SC SC SC ZT FB FA SC ZT FB FC SC SC SC SC	2 SC SC SC SC SC ZT FB FA SC FD FB FD SC SC SC SC	1 --- -- -- -- FF -- FH --- FH FH FG FG FG FG	2 FE --- -- FE FK -- FH --- FH FH FG FG FG FG	1 FL FM FM -- FM FM FM -- FM FO FO -- FM FO FO FO	2 FL FM FM -- FM FM FM -- FM FO FO -- FM FO FO FO
3 SC SC SC SC SC ZT FB FA SC FD FB FD SC SC SC SC	4 SC SC SC SC SC ZA FB FA SC ZT FB FD SC SC SC SC	3 SC SC SC SC SC ZT FB FA SC FD FB FD SC SC SC SC	4 SC SC SC SC SC ZA FB FA SC ZT FB FD SC SC SC SC	3 FE --- -- FE FI -- FH --- FH FH FG FG FG FG	4 --- -- -- FJ -- ZA FH FJ -- FH FH FG FG FG FG	3 FL FM -- -- FM FM FM -- FM FO ZA FO FM FO FO FO	4 FL FM FM -- FM FM FM -- FM -- FO FO FM FO FO FO
Vector Control Vector Length Window E		Instruction Parcel Translation F		Local Memory 32-47 G		Floating Multiply Segment Two 32-16 H	
1 --- -- -- --- AO AP AL AL --- VY VY VY VY	2 --- -- -- --- AN AP AL AL --- VY VY VY VY	1 GW GW GW GW BC --- -- BO --- BK BP BQ BP BQ	2 --- -- -- BB --- BN BB --- BL BR BS BR BS	1 SC SC SC SC WE WE WE WE WD WC WC SC SC SC SC SC	2 SC SC SC SC ZT WB WB ZT WD WC WC SC SC SC SC SC	1 FL FM FM -- FM FM FM -- FM FO FO -- FM FO FO FO	2 FL FM FM -- FM FM FM -- FM FO FO -- FM FO FO FO
3 --- -- -- --- AM -- VX VX VX VX AJ AJ AJ AJ	4 --- -- -- --- ZA AO -- VX VX VX VX AJ AJ AJ AJ	3 AQ AQ BZ BZ --- GX AR BT --- BY BU BV BW BX	4 --- BA BZ BZ --- -- -- --- ZA -- --- -- --	3 SC SC SC SC ZT WF WG ZT WD WC WC SC SC SC SC SC	4 SC SC SC SC WA WA ZA ZT WD WC WC SC SC SC SC SC	3 FL FM -- -- FM FM FM -- FM FO ZA FO FM FO FO FO	4 FL FM FM -- FM FM FM -- FM -- FO FO FM FO FO FO
Scalar Registers 32-35 Scalar Logical 32-35 Scalar Add 32-35 Vector Registers 32-35 I		Scalar Registers 36-39 Scalar Logical 36-39 Scalar Add 36-39 Vector Registers 36-39 J		Floating Multiply Segment One 16-00 K		Floating Multiply Segment Four Module Clock Distribution L	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 FM FM FM FM FM FM FL FL FO --- FO FS FO FO --	2 FM FM FM FM FM FM FL FM FO FO -- FO -- FO FO FO	1 FQ FQ FQ FQ -- ZA -- FN FQ FQ FQ FQ -- FP -- --	2 FQ FQ FQ FQ -- FQ -- WL FQ FQ FQ FQ -- FP -- --
3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY --- ZA -- -- GA VU GC	3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY --- ZA -- VU GA ---	3 FM FM FM FM FM FM FL FM FO -- ZA FO FO FO FO --	4 FM FM FM FM FM FM FL FM FO FO -- FO -- FO FO FO	3 FQ FQ FQ FQ --- ZA -- FQ FQ FQ FQ -- FP -- --	4 FQ FQ FQ FQ --- FQ -- WL FQ FQ FQ FQ -- FP -- --
Scalar Registers 40-43 Scalar Logical 40-43 Scalar Add 40-43 Vector Registers 40-43 Vector Logical 32-47 Vector Mask 32-47 Real Time Clock 32-47 M		Scalar Registers 44-47 Scalar Logical 44-47 Scalar Add 44-47 Vector Registers 44-47 Vector Add 32-47 Comp Iota Mask 32-47 N		Floating Multiply Exponent O		Floating Multiply Final Coefficient P	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 --- -- -- --- YD --- YA YB -- --- YC	2 --- -- -- --- YE --- YA YB -- FZ FZ -- YC	1 FR FS FU -- FY --- FU -- FX FX FX FZ --- --	2 -- FR FV -- -- ZA -- FV FX -- FX FX -- FZ FZ --
3 VS VS VS VS VJ VK VL AZ --- -- -- VT VT VT VT	4 VG VG VG VG VH VH VH VH VI -- ZA VI -- GA VU --	3 VS VS VS VS --- -- AZ --- ZA -- VT VT VT VT	4 VA VA VA VA VB VB VB VB VC VD VE -- -- GA VU --	3 --- -- -- --- YF -- --- YA YB -- FZ FZ -- YC	4 --- -- -- --- YG --- YA YB -- --- YC	3 FR -- FW FU --- FU -- FX FX -- FX -- FZ FZ --	4 -- FR FU FV FT -- FV -- -- FX FX -- FZ FZ FZ --

D Module

Background Processor

Instruction Issue Control A		Instruction Stack Buffer Registers B		Instruction Stack Buffer Control C		Foreground Channel Buffer D	
1 CC CN CM CM CP CP ZA -- CB CZ CY CW CA CD -- --	2 CG CN CF -- CP CP CM CM -- CZ CY CW -- CH CE --	1 VR VR VR VR VR VR VR VR VR VR VR VR VR VR VR VR	2 VR VR VR VR VR VR VR VR VR VR VR VR VR VR VR VR	1 -- -- -- -- VR VR -- -- PF PJ -- -- -- PG PG	2 -- PE -- -- -- -- -- -- PI PH -- -- -- PG PG	1 HR HR HU -- VV VV HP -- VV VV -- HS HQ HQ -- HW	2 HR HR HU -- VV VV HP -- VV VV -- HT HQ HQ -- HX
3 -- CN CM CM CP CP CS CT CI CZ CY CW CJ CK CR --	4 CO CO -- -- CP CP CM CM CX CZ CY CW -- CL CQ --	3 PB PB PB PB PB PB -- PB PB PB ZA PB PB PB PB PB	4 PC PC PC PC PA PA PB PA PA PA PB PA PC PC PC PC	3 -- PD -- -- -- -- -- -- -- ZA -- -- -- PG PG	4 -- -- -- -- -- -- -- -- -- -- -- -- -- PG PG	3 HR HR HU -- VV VV HP -- VV VV -- HV HQ HQ -- HY	4 HR HR HU -- VV VV HP -- VV VV ZA -- HQ HQ -- --
Vector Control Vector Length Window E		Instruction Parcel Translation F		Local Memory G		Foreground Channel Interface H	
						Module Clock Distribution	
1 -- -- AP AP -- -- AO AP AL AL -- -- VY VY VY VY	2 -- -- AP AP -- -- AN AP AL AL -- -- VY VY VY VY	1 GS GO GQ GQ BC GO -- GP BO GO -- BK BP BQ BP BQ	2 -- -- -- -- BB -- -- BN BB -- -- GR BR BS BR BS	1 SC SC SC SC WE WE WE WE WD WC WC SC SC SC SC SC	2 SC SC SC SC ZT WB WB ZT WD WC WC SC SC SC SC SC	1 HK HK -- HB HE HI -- HC -- -- -- -- -- -- HD --	2 HK HK -- HB HF HJ -- HL -- -- -- HM -- -- -- HN
3 -- -- -- -- -- AM -- -- VX VX VX VX AJ AJ AJ AJ	4 -- -- -- -- -- ZA AO -- VX VX VX VX AJ AJ AJ AJ	3 AQ AQ BZ BZ -- -- -- AR BT -- -- BY BU BV BW BX	4 -- BA BZ BZ -- -- -- -- -- ZA -- -- -- -- -- --	3 SC SC SC SC ZT WF WG ZT WD WC WC SC SC SC SC SC	4 SC SC SC SC WA WA ZA ZT WD WC WC SC SC SC SC SC	3 HK HK -- HB HG HI -- -- -- ZA -- -- -- -- -- --	4 HK HK -- HB HH HI -- HA -- -- ZA -- -- -- -- --
Scalar Registers 48-51 Scalar Logical 48-51 Scalar Add 48-51 Vector Registers 48-51		Scalar Registers 52-55 Scalar Logical 52-55 Scalar Add 52-55 Vector Registers 52-55		Common Memory Read Packets 0, 1, 2		Common Memory Read Packet 3 Data Merge	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 XZ VV XN XK VV XN XK XQ VV XN XK XQ XZ VV XN XK	2 XZ VV XN XK VV XN XK XQ VV XN XK XQ XZ VV XN XK	1 -- -- XI XI VV XV XI XI XZ XP XI XO VV XP XI XO	2 -- -- XI XI VV XV XI XI XZ XK XI XO VV XK XI XO
3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY -- -- ZA -- -- GA VU GD	3 VS VS VS VS AT AU AV AV AS AS AS AS VT VT VT VT	4 AS AS AS AS AW AX AX AY -- -- ZA -- VU GA -- --	3 XZ VV XN XK -- -- XK XQ -- -- XK XQ XZ VV XN XK	4 -- VV XN -- -- XW ZA XQ -- XW XV XQ -- VV XN --	3 -- -- -- -- XL XL ZA -- XL XL XI XI XX -- XI XI	4 -- -- -- -- XL XL XJ XM XL XL XJ XM XQ XQ XQ XM
Scalar Registers 56-59 Scalar Logical 56-59 Scalar Add 56-59 Vector Registers 56-59 Vector Logical 48-63 Vector Mask 48-63 Real Time Clock 48-63		Scalar Registers 60-63 Scalar Logical 60-63 Scalar Add 60-63 Vector Registers 60-63 Vector Add 48-63 Comp Iota Mask 48-63		Vector Shift Vector Population Count Vector Leading Zero Count		Common Memory Read SEC-DED	
1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	2 VQ VQ VQ VQ VR VR VR VR VP VP VP VP VT VT VT VT	1 UG UA UE UG UH -- -- UH UI UB UF UI UC UD -- UL	2 UG UA UE UG UH -- -- UH UI -- -- UI UL UJ -- UL	1 XP VV XZ VV XP XQ XI XI -- XR XS -- XU XU XU XU	2 XP VV XZ VV XP XQ XI XI -- XR XS -- XU XU XU XU
3 VS VS VS VS VJ VK VL AZ -- -- -- -- VT VT VT VT	4 VG VG VG VG VH VH VH VH VI -- ZA VI -- GA VU --	3 VS VS VS VS -- -- -- AZ -- -- ZA -- VT VT VT VT	4 VA VA VA VA VB VB VB VB VC VD VE VF -- GA VU --	3 UG UA UE UG UH -- ZA UH UI -- -- UI UL UK -- UL	4 UG UA UE UG UH -- -- UH UI -- -- UI UL -- -- UL	3 XP VV XZ XW XP XN XO XN XV XR XT -- XU XU XU XU	4 XP XQ XY XQ XP XN XO XN -- XR ZA -- XU XU XU XU

E Module

Foreground Processor

A		B		Local Memory C		Foreground Interface To Common Memory (Test Circuits Only) D	
1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 ----- SC SC SC SC -- ZT -- LA LA LA LA	2 ----- SC SC SC SC -- ZT -- LB LB LB LB	1 -- NY -- -- NS NX NV NX -- NZ NW -----	2 -- NY -- -- NT NX NV NX -- NZ NW -----
3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 ----- SC SC SC SC -- ZT -- LC LC LC LC	4 ----- SC SC SC SC SC ZT ZA SC -----	3 ----- -- NU NX NV NX -- NZ NW -----	4 ----- -- NX NV NX ZA NZ NW -----
E		Instruction Memory 0-32 K F		A, B, C Registers Add Shift Logical Real Time Clock G		Channel Loop Zero Channel Loop One Channel Loop Two Channel Loop Three H	
1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 SC SC SC SC SC SC SC SC SC -- ZT ME -- MC MC MD	2 SC SC SC SC SC SC SC SC SC ZT ZT ME -- MC MC --	1 LD LD LD LD -- LE LE -- MS MS MS MS MR MR MR MR	2 LD LD LD LD ----- MS MS MS MS MR MR MR MR	1 -- -- LK LK -- LJ -- -- LJ -- -- LM LK LK	2 -- -- LK LK -- LJ -- -- LJ -- -- LM LK LK
3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 SC SC SC SC SC SC SC SC SC -- ZT ME -- MC MC MF	4 SC SC SC SC SC SC SC SC SC ZA ZT ME -- MC MC --	3 LG LG LG LG LI LI -- -- ZA -- MR MR MR MR	4 LF LF LF LF LH LH LH LH ----- MR MR MR MR	3 -- -- LK LK LL -- -- LL -- -- -- -- LK LK	4 -- -- LK LK LL -- -- LL -- ZA -- -- -- LK LK
I		Instruction Memory 32-64 K J		Console Interface Instruction Issue Control Instruction Memory Address Control Data Assembly/Disassembly K		System Deadstart Distribution Clear Real Time Clock Distribution Module Clock Distribution L	
1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 SC SC SC SC SC ZT ME ME -- MC MC -- SC SC SC SC	2 SC SC SC SC SC ZT ME ME -- MC MC ZT SC SC SC SC	1 MJ MI MU -- MA MN MP -- MG MH -- ZT --	2 MK MI MV -- MA MT MZ -- MG MH -- -----	1 ----- ----- ----- -----	2 ----- ----- ----- -----
3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 SC SC SC SC SC ZT -- -- MC MC MF SC SC SC SC	4 SC SC SC SC SC ZT ZA -- -- MC MC MD SC SC SC SC	3 ML MI MY -- MA MM MW MX MG MH -- -----	4 MB MI -- MQ MB MO ZA -- MG MH -- LO -- LN --	3 ----- ----- ----- -----	4 ----- ----- -- ZA ZA -- -----
M		N		O		System Clock Distribution P	
1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 ----- ----- ----- -----	2 ----- ----- ----- -----	1 ----- ----- ----- -----	2 ----- ----- ----- -----
3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 ----- ----- ----- -----	4 ----- -- ZA -- ----- -----	3 ----- ----- -- ZA -- -----	4 ----- ----- -- ZA -- -----

Common Memory

Common Memory Bank 0 Bits 44-53 Check Bits E and F		Common Memory Bank 0 Bits 22-31 Check Bits C and D		Common Memory Bank 0 Bits 00-11		Common Memory Bank 0 Control Bank Pointer Distribution to G Modules	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EJ EJ EJ TG ¹	--- -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ EJ TG	EQ -- ZA TK
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EK EK EK --	EQ EM -- TK
						-- -- -- TF	TE TE -- --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	--- -- -- ³	--- -- -- ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	---	---
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	---	---
						---	---
Common Memory Bank 0 Bits 54-63 Check Bits G and H		Common Memory Bank 0 Bits 32-43		Common Memory Bank 0 Bits 12-21 Check Bits A and B		Common Memory Bank Reservation Flags Bank 0 and 1 Address Registers	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI QD ¹	EI EI -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QD QD	EI EI ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QC QC	EI EI QC QC
						EI EI QC QC	EI EI QC QC
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	EI QB QB QB ³	-- QB QB QB ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI QB QB QB	-- QB QB QB
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI QB QC QC	-- QB QC QC
						EI QB QC QC	-- QB QC QC
Common Memory Bank 1 Bits 54-63 Check Bits G and H		Common Memory Bank 1 Bits 32-43		Common Memory Bank 1 Bits 12-21 Check Bits A and B		Common Memory Go Bank Generation Bank 0 and 1 Address Registers	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI QG QG ¹	EI EI QG QG ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QH --	EI EI QH QS
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QE QI	EI EI QE QS
						EI EI -- QI	EI EI -- QF
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	ES QB QG QG ³	-- -- QG QG ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER QB QH QA	-- -- ZA QA
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER QB QE QA	-- -- QE QA
						-- QB -- --	-- -- -- --
Common Memory Bank 1 Bits 44-53 Check Bits E and F		Common Memory Bank 1 Bits 22-31 Check Bits C and D		Common Memory Bank 1 Bits 00-11		Common Memory Bank 1 Control Bank Pointer Distribution to G Modules Module Clock Distribution Octant Clock Distribution	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	--- -- -- ¹	TE TE -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EK EK EK --	EQ EM ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ EJ TG	EQ -- -- --
						EJ EJ EJ TG	-- -- -- --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	--- -- -- ³	--- -- -- ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	---	---
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	---	---
						---	---

G Module

Common Memory

Common Memory Bank 0 Bits 44-53 Check Bits E and F		Common Memory Bank 0 Bits 22-31 Check Bits C and D		Common Memory Bank 0 Bits 00-11		Common Memory Bank 0 Control Bits 0, 1 Write Data Packet to Octant Banks 0	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EJ EJ -- -- ¹	EK EK EK -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ -- --	-- -- ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ QN QN	-- EM QN QN
						TE QL QN QN	TE QL QN QN
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	QM QM QM QM ³	QM QM QM QM ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	QM QM QM QM	QM QM QM QM
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EQ -- QN QN	-- -- QN QN
						EQ QL QN QN	-- QL QN QN
Common Memory Bank 0 Bits 54-63 Check Bits G and H		Common Memory Bank 0 Bits 32-43		Common Memory Bank 0 Bits 12-21 Check Bits A and B		Common Memory Write Data Merge Bank 0 and 1 Address Registers Bit 0, 1 Write Data from Background Priority A, B	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI -- -- ¹	EI EI -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QK QK	EI EI QK QK
						EI EI TJ TJ	EI EI TJ TJ
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	EI -- -- -- ³	-- -- QP QP ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI -- -- --	-- -- -- --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI -- QK QK	-- -- QK QK
						EI -- TJ TJ	-- -- -- --
Common Memory Bank 1 Bits 54-63 Check Bits G and H		Common Memory Bank 1 Bits 32-43		Common Memory Bank 1 Bits 12-21 Check Bits A and B		Common Memory Write Data Merge Bank 0 and 1 Address Registers Bits 0, 1 Write Data from Background Priority C, D Module Clock Distribution	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI TJ TJ ¹	EI EI TJ TJ ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI QK QK	EI EI QK QK
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI -- --
						EI EI -- --	EI EI -- --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	ES -- TJ TJ ³	-- -- -- -- ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER -- QK QK	-- -- QK QK
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER -- ZA --	-- -- ZA --
						-- -- -- --	-- -- QP QP
Common Memory Bank 1 Bits 44-53 Check Bits E and F		Common Memory Bank 1 Bits 22-31 Check Bits C and D		Common Memory Bank 1 Bits 00-11		Common Memory Bank 1 Control Bits 0, 1 Write Data Packet to Octant Banks 1	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	TE QL QN QN ¹	TE QL QN QN ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ QN QN	-- EM QN QN
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ -- --	-- -- ZA --
						EJ EJ -- --	EK EK EK --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	EQ QL QN QN ³	-- QL QN QN ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EQ -- QN QN	-- -- QN QN
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	QM QM QM QM	QM QM QM QM
						QM QM QM QM	QM QM QM QM

H Module

Common Memory

Common Memory Bank 0 Bits 44-53 Check Bits E and F		Common Memory Bank 0 Bits 22-31 Check Bits C and D		Common Memory Bank 0 Bits 00-11		Common Memory Bank 0 Control Common Memory Readout Data Merge	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EJ EJ EJ -- ¹	-- -- -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ EJ --	EQ -- ZA TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EK EK EK --	EQ EM -- TD
						-- -- QQ QQ	TE TE QQ QQ
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	-- -- -- -- ³	-- -- -- -- ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- -- TD	-- -- -- TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- -- TD	-- -- -- TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- QQ QQ	-- -- QQ QQ
Common Memory Bank 0 Bits 54-63 Check Bits G and H		Common Memory Bank 0 Bits 32-43		Common Memory Bank 1 Bits 12-21 Check Bits A and B		Common Memory Readout Data Merge Bank 0 and 1 Address Registers Bits 0, 1, 2 Group Address to Banks 0 and 1	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI -- -- ¹	EI EI -- -- ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI -- --
						EI EI QW TC	EI EI QW TC
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	EI -- -- -- ³	-- -- -- -- ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI -- -- --	-- -- -- --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI -- -- --	-- -- -- --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI -- QW TC	-- -- QW QX
Common Memory Bank 1 Bits 54-63 Check Bits G and H		Common Memory Bank 1 Bits 32-43		Common Memory Bank 1 Bits 12-21 Check Bits A and B		Common Memory Readout Data Merge Bank 0 and 1 Address Registers Final Readout Merge to Background Processors Module Clock Distribution	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	EI EI -- QR ¹	EI EI -- QR ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI -- --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EI EI -- --	EI EI -- --
						EI EI -- --	EI EI -- --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	ES -- -- QR ³	-- -- -- QR ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER -- -- --	-- -- ZA --
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	ER -- -- --	-- -- ZA --
						-- -- -- --	-- -- -- --
Common Memory Bank 1 Bits 44-53 Check Bits E and F		Common Memory Bank 1 Bits 22-31 Check Bits C and D		Common Memory Bank 1 Bits 00-11		Common Memory Bank 1 Control Common Memory Readout Data Merge	
SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	SD SD SD SD ¹	SD SD SD SD ²	-- -- QQ QQ ¹	TE TE QQ QQ ²
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EK EK EK --	EQ EM -- TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	EJ EJ EJ --	EQ -- -- TD
						EJ EJ EJ --	-- -- -- --
SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	SD SD SD SD ³	SD SD SD SD ⁴	-- -- QQ QQ ³	-- -- QQ QQ ⁴
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- -- TD	-- -- ZA TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- -- TD	-- -- -- TD
SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	SD SD SD SD	-- -- -- --	-- -- -- --

K Module

I/O

DD49, DS40, and LSX

Note: K Module
is no longer
used.

DD-49 Zero Buffer A A		DD-49 Zero Buffer B B		Node Address Distribution C		DD-49 or DS-40 Disk Controller Zero D	
SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	--- 1 --- --- 2 --- --- 3 --- --- 4 ---	--- WL --- --- 2 --- --- 3 --- --- 4 ---	--- 1 --- WQ WQ -- WO --- WZ WY WV -- WT --	--- 2 --- WP WP -- WO --- ZT WR -- WV WS WT --
SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	--- WL --- --- 3 --- --- 4 --- --- 5 ---	--- 4 --- --- ZA --- --- WX --- --- WT ---	--- 3 --- WP WP -- WO --- WX -- WV WU WT --	--- 4 --- WQ WQ ZA WO ZT -- -- WN WV -- WT --
DD-49 One Buffer A E		DD-49 One Buffer B F		LSX Input Buffer G		DD-49 or DS-40 Disk Controller One H Module Clock Distribution	
SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	--- 1 --- WQ WQ -- WO --- WZ WY WV -- WT --	--- 2 --- WP WP -- WO --- ZT WR -- WV WS WT --
SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	--- 3 --- WP WP -- WO --- ZA WX -- WV WU WT --	--- 4 --- WQ WQ ZA WO ZT -- -- WN WV -- WT --
DD-49 Two Buffer A I		DD-49 Two Buffer B J		Low-Speed External I/O Controller K		DD-49 or DS-40 Disk Controller Two L	
SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	--- 1 --- ND ND -- NB --- NH -- NE NL NK NN --	--- 2 --- NC NC -- NB --- ZT -- NF NL -- NN --	--- 1 --- WQ WQ -- WO --- WZ WY WV -- WT --	--- 2 --- WP WP -- WO --- ZT WR -- WV WS WT --
SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	--- 3 --- NC NC -- NB ZT -- -- NG NL NJ NN --	--- 4 --- ND ND ZA NB --- NH -- NA NL NM NN --	--- 3 --- WP WP -- WO --- WX -- WV WU WT --	--- 4 --- WQ WQ ZA WO ZT -- -- WN WV -- WT --
DD-49 Three Buffer A M		DD-49 Three Buffer B N		LSX Output Buffer O		DD-49 or DS-40 Disk Controller Three P	
SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC	SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC	--- 1 --- WQ WQ -- WO --- WZ WY WV -- WT --	--- 2 --- WP WP -- WO --- ZT WR -- WV WS WT --
SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC	SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC	--- 3 --- WP WP -- WO --- WX -- WV WU WT --	--- 4 --- WQ WQ ZA WO ZT -- -- WN WV -- WT --

L Module

I/O

HIPPI Channels

Four 32-Bit HIPPI Channel Controller Pairs

HIPPI Destination Buffer Zero		HIPPI Source Buffer Zero		HIPPI Destination Controller Zero		HIPPI Source Controller Zero	
1	2	1	2	1	2	1	2
SC SC SC WI	SC SC SC WH	SC SC SC WI	SC SC SC WH	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT ZT	SC WC ZT WH	SC WC ZT ZT	SC WC ZT WH	-- JG JE JB	JD JO JF JB	-- KG KE KB	KD KH KF KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR JI JA	KD VR -- --	-- VR -- KA
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JL JZ	JK JK JL --	KK KK -- KZ	KK KK KL --
3	4	3	4	3	4	3	4
SC SC SC WK	SC SC SC WJ	SC SC SC WK	SC SC SC WJ	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT WK	SC WC ZA ZT	SC WC ZT WK	SC WC ZA ZT	-- JJ JP JB	JD JO JQ JB	-- KJ -- KB	KD -- -- KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR ZA --	KD VR -- KI	KO VR ZA --
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JM JZ	JK JK JN --	KK KK KM KP	KK KK KN KZ
HIPPI Destination Buffer One		HIPPI Source Buffer One		HIPPI Destination Controller One		HIPPI Source Controller One	
1	2	1	2	1	2	1	2
SC SC SC WI	SC SC SC WH	SC SC SC WI	SC SC SC WH	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT ZT	SC WC ZT WH	SC WC ZT ZT	SC WC ZT WH	-- JG JE JB	JD JO JF JB	-- KG KE KB	KD KH KF KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR JI JA	KD VR -- --	-- VR -- KA
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JL JZ	JK JK JL --	KK KK -- KZ	KK KK KL --
3	4	3	4	3	4	3	4
SC SC SC WK	SC SC SC WJ	SC SC SC WK	SC SC SC WJ	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT WK	SC WC ZA ZT	SC WC ZT WK	SC WC ZA ZT	-- JJ JP JB	JD JO JQ JB	-- KJ -- KB	KD -- -- KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR ZA --	KD VR -- KI	KO VR ZA --
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JM JZ	JK JK JN --	KK KK KM KP	KK KK KN KZ
HIPPI Destination Buffer Two		HIPPI Source Buffer Two		HIPPI Destination Controller Two		HIPPI Source Controller Two	
1	2	1	2	1	2	1	2
SC SC SC WI	SC SC SC WH	SC SC SC WI	SC SC SC WH	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT ZT	SC WC ZT WH	SC WC ZT ZT	SC WC ZT WH	-- JG JE JB	JD JO JF JB	-- KG KE KB	KD KH KF KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR JI JA	KD VR -- --	-- VR -- KA
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JL JZ	JK JK JL --	KK KK -- KZ	KK KK KL --
3	4	3	4	3	4	3	4
SC SC SC WK	SC SC SC WJ	SC SC SC WK	SC SC SC WJ	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT WK	SC WC ZA ZT	SC WC ZT WK	SC WC ZA ZT	-- JJ JP JB	JD JO JQ JB	-- KJ -- KB	KD -- -- KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR ZA --	KD VR -- KI	KO VR ZA --
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JM JZ	JK JK JN --	KK KK KM KP	KK KK KN KZ
HIPPI Destination Buffer Three		HIPPI Source Buffer Three		HIPPI Destination Controller Three		HIPPI Source Controller Three	
1	2	1	2	1	2	1	2
SC SC SC WI	SC SC SC WH	SC SC SC WI	SC SC SC WH	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT ZT	SC WC ZT WH	SC WC ZT ZT	SC WC ZT WH	-- JG JE JB	JD JO JF JB	-- KG KE KB	KD KH KF KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR JI JA	KD VR -- --	-- VR -- KA
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JL JZ	JK JK JL --	KK KK -- KZ	KK KK KL --
3	4	3	4	3	4	3	4
SC SC SC WK	SC SC SC WJ	SC SC SC WK	SC SC SC WJ	JC -- VR --	-- JC VR --	KC -- WL --	-- KC -- --
SC WC ZT WK	SC WC ZA ZT	SC WC ZT WK	SC WC ZA ZT	-- JJ JP JB	JD JO JQ JB	-- KJ -- KB	KD -- -- KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR ZA --	KD VR -- KI	KO VR ZA --
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JM JZ	JK JK JN --	KK KK KM KP	KK KK KN KZ

M Module

I/O

DD49, DS40, LSX and HIPPI Channels

HIPPI Destination Buffer Zero		HIPPI Source Buffer Zero		HIPPI Destination Controller Zero		HIPPI Source Controller Zero	
1 SC SC SC WI	2 SC SC SC WH	1 SC SC SC WI	2 SC SC SC WH	1 JC -- VR --	2 -- JC VR --	1 KC -- WL --	2 -- KC -- --
SC WC ZT ZT	SC WC ZT WH	SC WC ZT ZT	SC WC ZT WH	-- JG JE JB	JD JO JF JB	-- KG KE KB	KD KH KF KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR JI JA	KD VR -- --	-- VR -- KA
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JL JZ	JK JK JL --	KK KK -- KZ	KK KK KL --
3 SC SC SC WK	4 SC SC SC WJ	3 SC SC SC WK	4 SC SC SC WJ	3 JC -- VR --	4 -- JC VR --	3 KC -- WL --	4 -- KC -- --
SC WC ZT WK	SC WC ZA ZT	SC WC ZT WK	SC WC ZA ZT	-- JJ JP JB	JD JO JQ JB	-- KJ -- KB	KD -- -- KB
SC WC WD WE	SC WC WD WE	SC WC WD WE	SC WC WD WE	JD VR JI --	-- VR ZA --	KD VR -- KI	KO VR ZA --
SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	JK JK JM JZ	JK JK JN --	KK KK KM KP	KK KK KN KZ
E		DD-49 One Buffer		LSX Input Buffer		DD-49 or DS-40 Disk Controller One	
Module Clock Distribution		F		G		H	
1 -- -- -- --	2 -- -- -- --	1 SC SC SC WI	2 SC SC SC WH	1 SC SC SC SC	2 SC SC SC SC	1 -- TS UM --	2 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT ZT	SC WC ZT WH	ZT WE NP ZT	ZT WE NQ ZT	TR -- -- TM	TR TP TQ TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	WD NO -- SC	WD NO -- SC	TS ZA -- --	TS ZT -- --
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	UN TT TW --	TZ TU TW --
3 -- -- -- --	4 -- -- -- --	3 SC SC SC WK	4 SC SC SC WJ	3 SC SC SC SC	4 SC SC SC SC	3 -- TS UM --	4 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT WK	SC WC ZA ZT	ZT WE NR ZT	ZT WE ZA ZT	TR -- -- TM	TR -- ZA TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	WD NO -- SC	WD NO -- SC	TS TV TO TN	TS ZT -- TL
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	UN TY TW --	TZ TX TW --
I		DD-49 Two Buffer		LSX Sequence Control		DD-49 or DS-40 Disk Controller Two	
1 -- -- -- --	2 -- -- -- --	1 SC SC SC WI	2 SC SC SC WH	1 NL NF -- NN	2 NL -- -- NE	1 -- TS UM --	2 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT ZT	SC WC ZT WH	ZT -- -- NA	ZT -- -- NC	TR -- -- TM	TR TP TQ TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	-- NJ -- NB	-- NI -- NB	TS -- -- --	TS ZT -- --
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	-- -- -- NN	-- -- -- --	UN TT TW --	TZ TU TW --
3 -- -- -- --	4 -- -- -- --	3 SC SC SC WK	4 SC SC SC WJ	3 NL -- -- NH	4 NL -- -- NH	3 -- TS UM --	4 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT WK	SC WC ZA ZT	ZT -- -- NG	ZT -- ZA --	TR -- -- TM	TR -- ZA TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	-- NK -- NB	-- -- -- NB	TS TV TO TN	TS ZT -- TL
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	ND -- -- NM	ND -- -- NM	UN TY TW --	TZ TX TW --
M		DD-49 Three Buffer		LSX Output Buffer		DD-49 or DS-40 Disk Controller Three	
1 -- -- -- --	2 -- -- -- --	1 SC SC SC WI	2 SC SC SC WH	1 SC SC SC SC	2 SC SC SC SC	1 -- TS UM --	2 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT ZT	SC WC ZT WH	ZT WE NP ZT	ZT WE NQ ZT	TR -- -- TM	TR TP TQ TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	WD NO -- SC	WD NO -- SC	TS -- -- --	TS ZT -- --
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	UN TT TW --	TZ TU TW --
3 -- -- -- --	4 -- -- -- --	3 SC SC SC WK	4 SC SC SC WJ	3 SC SC SC SC	4 SC SC SC SC	3 -- TS UM --	4 -- TS UM --
-- -- -- --	-- -- -- --	SC WC ZT WK	SC WC ZA ZT	ZT WE NR ZT	ZT WE ZA ZT	TR -- -- TM	TR -- ZA TM
-- -- -- --	-- -- -- --	SC WC WD WE	SC WC WD WE	WD NO -- SC	WD NO -- SC	TS TV TO TN	TS ZT -- TL
-- -- -- --	-- -- -- --	SC SC SC SC	SC SC SC SC	SC SC SC SC	SC SC SC SC	UN TY TW --	TZ TX TW --

N Module

I/O

HIPPI

Dynamically Switchable Four 32-Bit Controller Pairs or

Two 64-Bit Controller Pairs or

Two 32-Bit and One 64-Bit Controller Pairs

HIPPI Destination Buffer Zero				HIPPI Source Buffer Zero				Controller Design in Process for the N Module as of July 28, 1993.							
A				B				C				D			
1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC				1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC			
3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC				3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC			
HIPPI Destination Buffer One				HIPPI Source Buffer One				G				H			
1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC				1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC			
3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC				3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC			
HIPPI Destination Buffer Two				HIPPI Source Buffer Two				K				L			
1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC				1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC			
3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC				3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC			
HIPPI Destination Buffer Three				HIPPI Source Buffer Three				O				P			
1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC				1 SC SC SC WI SC WC ZT ZT SC WC WD WE SC SC SC SC				2 SC SC SC WH SC WC ZT WH SC WC WD WE SC SC SC SC			
3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC				3 SC SC SC WK SC WC ZT WK SC WC WD WE SC SC SC SC				4 SC SC SC WJ SC WC ZA ZT SC WC WD WE SC SC SC SC			

List of Figures

Figures In Order of Appearance

- Four-Processor System Cabinet and Control Pod. 2
Eight-Processor System Cabinet. 3
Module Placement in the CRAY-3. 5
Module Types Used in a CRAY-3 Octant. 6
CRAY-3 Electronic Module. 7
Layers in a Module. 8
Exploded View of Layers in a Module. 9
Flex Connector. 10
Sizes of the Basic Components in a CRAY-3 Module.
11
Basic Components of the CRAY-3. 14
Common Memory Stack Usage. 17
Block Diagram of a Background Processor. 20
Background Processor Functional Divisions Among the
Modules. 22
Data Channels (One of Four). 23
Block Diagram of the Foreground Processor. 24
4.0 Micron Transistor—Semiconducting Channel. 31
4.0 Micron Transistor—N+ Implants. 32
4.0 Micron Transistor—Ohmic Metal. 32
4.0 Micron Transistor—First Layer Metal. 33
4.0 Micron Transistor—Recess Opening. 33
Electrical Characteristics of a Transistor. 34
4.0 Micron Current Source—First Layer Metal. 35
Type A Diode (3 Square Micron)—N- and N+ Implants.
36
Type A Diode—Metal Layers and Recess Insulator. 37
Design of a 4 Square Micron Type B Diode. 38
Diode Electrical Characteristics. 39
Design of a 95 Square Type C Diode/Capacitor. 40
Type C Diode/Capacitor—Ohmic Metal. 41
Type C Diode/Capacitor—First Layer Metal. 42
Type C Diode/Capacitor—Recess Insulator. 43
Capacitor Broken Into Two Pieces. 44
Capacitor Electrical Characteristics. 45
Design of a 700 Ohm Resistor—N- and N+ Implants. 46
Ohmic and First Layer Metal in a 700 Ohm Resistor. 47
Logic Package Block Diagrams. 49
Logic Package Power Grid and Cell Positions. 50
Logic Package Pinout. 51
DD Cell Schematic Diagram. 53
ED Cell Circuit Schematic. 55
FE Cell Circuit Schematic. 56
FQ Cell Circuit Schematic. 57
FF Cell Circuit Schematic. 58
FJ Cell Circuit Schematic. 59
CA Cell Circuit Schematic. 61
CC Cell Circuit Schematic. 63
AF Cell Circuit Schematic. 64
AG Cell Circuit Schematic. 65
AH Cell and AI Cell Circuit Schematics. 66
Vector Register Block Diagram. 67
VR Write Timing. 68
VR Read Timing. 69
Size Comparisons of the Memory Die. 70
SC Die Pinout. 71
SD and SE Die Pinouts. 72
SM Die Pinouts. 73

Generic IC Assembly. 80
 Parts of a Macro Text Line. 82
 Macros and the Generic IC Assembly. 83
 First Part of the AA Macro. 84
 Placement of the AB Macro. 85
 AC Macro Calls. 85
 Placement of the AC Macros. 86
 The Macro Lines for the Bonding Targets. 87
 Creating a CRAY-3 IC. 88
 Step One: The Generic IC Assembly is Created. 89
 Step Two: Placing the Required Logic Cells. 90
 Cell Placement in the IC Boolean Document. 91
 Cell Placement in EI Package without Null Cells. 92
 Close-Up of Cells on West F Street. 93
 Step Three: Routing the Final IC Assembly. 94
 Bottom Left Corner of the EI Package. 95
 Routed Portion of C Street West and Adjacent Avenue. 96
 16-Up Mask Frame Layout. 103
 Data Rotation During Fracturing. 105
 Orientation of the FET Gates on a Wafer. 106
 Fractured Stripes and Segments. 107
 Fractured Data File Name Conventions. 108
 Understanding the Layer Title. 110
 Correlation of IC Positioning. 111
 Stepping Sequence on a GaAs Wafer. 116
 Stepped Die Patterns on a Wafer. 117
 Dielectric Encapsulation and N- Implant. 119
 Heavy Ion (N+) Implant. 119
 Ohmic Contact Metallization and Proton Isolation. 120
 Schottky Barrier and First Layer Metal Interconnect. 121
 Second Level Metal. 122
 CERMET (Optional) and Passivation. 122
 Wafer Dicing. 127
 Notcher Mechanism. 129
 Die Bonding Illustration. 130
 Gold Leads Bonded to a Portion of a Die. 131
 Exploded View of a CRAY-3 Module. 134
 Cross Sectional Diagram of a Module. 135
 Arrangement of Stacks and ICs. 136
 The Board Stack on a Module. 138
 Plated Hole in a Logic Circuit Board. 142
 P Stack Logic Circuit Board. 144
 Corner of a Circuit Board with Trace Lines. 145
 256K Memory Circuit Board. 148
 One Meg Memory Circuit Board. 149
 Module Logic Plate. 152
 Portion of a Logic Plate with Trace Lines. 153
 256 K Memory Resistor Plate. 156
 Resistors on a Logic Module Resistor Plate. 157
 256K Memory Module Power Plate. 159
 SM Die and Bypass Capacitors. 161
 Module Power Blade. 162
 Vp Power Distribution in a Module. 163
 Power Distribution in Logic Modules. 164
 Power Distribution in 256K Memory Modules. 165
 Power Distribution in Four Meg Memory Modules. 166
 Logic Connectors Connected to a Corner Logic Board. 167
 Module Logic Connector with Flex Circuit. 169
 Module Build Flow Chart. 172
 The Kitting of Die for a Module. 173
 Small Wires Soldered to Circuit Boards. 174
 Flex Connector Re-Flow Fixture. 176
 Flex Circuits Attached to Board Wirebonds. 177
 Die Attach Stamping Fixture. 179
 Twist Pin Illustrations. 180
 Twist Pin Defects. 182
 Stacking a Module. 184
 Illustration of Stacking Layers. 185
 Twist Pin Seating. 187
 Pulse-Power Test Station. 191
 Pulse-Power Supply Block Diagram. 193
 Time Chart for Pulse-Power Testing. 194
 Block Diagram of the System Test Interface. 198
 Front View of a Four-Processor Tank. 208
 Front View of Eight-Processor System Cabinet. 209
 Eight-Processor System Cabinet with Skins. 210
 Front View of Prototype One- or Two-Processor System Cabinet. 211
 Sub-Floor Mounting of the System Cabinet. 212
 Baseplate for the Octants. 213
 Top Plate for a Four-Processor CRAY-3. 214
 Top Plate for an Eight-Processor CRAY-3. 215
 Modules Mounted in an Eight-Processor Cabinet. 216
 Top Plate for a Prototype One- or Two-Processor CRAY-3. 217
 Modules and Power Supplies in an Octant. 218
 Power Distribution in an Octant Assembly. 219
 Horizontal Power Buses for Memory from Module Side. 220
 Horizontal Power Buses for the Memory Modules. 220
 Horizontal Power Buses for Logic Modules from Module Side. 221
 Horizontal Power Buses for Logic Modules. 221
 Linear Power Supplies in an Octant Assembly. 223
 Simplified Top View of a Vp Logic Linear Power Supply. 224
 Simplified Top View of a Vp Memory Power Supply. 225
 Simplified Top View of a Vm Power Supply. 226
 Simplified Top View of a Vn Power Supply. 227
 Wire Harness Connection Points to the Modules. 229
 Simplified Diagram of Four-Processor Wire Harness. 230
 Simplified Diagram of the Cooling System. 232
 C-Pod Cooling Components—Front Right View. 233
 C-Pod Cooling Components—Front Left View. 234
 Translator Box. 236
 C-Pod External View. 238
 C-Pod-Front Right View. 239
 C-Pod Front Left View. 240

Positioning of User Controls and Maintenance Panel. 241
 C-Pod User Controls. 242
 The Touch Screen Display. 243
 Maintenance Panel Controls. 244
 Electronic Assemblies in the C-Pod—Front Left View. 246
 Electronic Components in the C-Pod—Front Right View. 249
 Sub-Floor Mounting of the C-Pod. 250
 Block Diagram of the DD-49/DS-40 Disk Controller. 253
 Drive Function Timing Sequence. 254
 Cable A Timing Diagram. 255
 Cable B Timing. 256

Alphabetical Listing of Figures

A

AC Macro Calls. 85
 AF Cell Circuit Schematic. 64
 AG Cell Circuit Schematic. 65
 AH Cell and AI Cell Circuit Schematics. 66
 Arrangement of Stacks and ICs. 136

B

Background Processor Functional Divisions Among the Modules. 22
 Baseplate for the Octants. 213
 Basic Components of the CRAY-3. 14
 Block Diagram of a Background Processor. 20
 Block Diagram of the DD-49/DS-40 Disk Controller. 253
 Block Diagram of the Foreground Processor. 24
 Block Diagram of the System Test Interface. 198
 The Board Stack on a Module. 138

C

C-Pod Cooling Components—Front Right View. 233
 C-Pod Cooling Components—Front Left View. 234
 Translator Box. 236
 C-Pod External View. 238
 C-Pod-Front Right View. 239
 C-Pod Front Left View. 240
 C-Pod User Controls. 242
 Cable A Timing Diagram. 255
 Cable B Timing. 256
 CA Cell Circuit Schematic. 61
 Capacitor Broken Into Two Pieces. 44
 Capacitor Electrical Characteristics. 45
 CC Cell Circuit Schematic. 63

Cell Placement in the IC Boolean Document. 91
 Cell Placement in EI Package without Null Cells. 92
 CERMET (Optional) and Passivation. 122
 Close-Up of Cells on West F Street. 93
 Common Memory Stack Usage. 17
 Simplified Diagram of the Cooling System. 232
 Corner of a Circuit Board with Trace Lines. 145
 Correlation of IC Positioning. 111
 CRAY-3 Electronic Module. 7
 Creating a CRAY-3 IC. 88
 Cross Sectional Diagram of a Module. 135

D

Data Channels (One of Four). 23
 DD Cell Schematic Diagram. 53
 Design of a 4 Square Micron Type B Diode. 38
 Design of a 95 Square Type C Diode/Capacitor. 40
 Design of a 700 Ohm Resistor—N- and N+ Implants. 46
 Diode Electrical Characteristics. 39
 Drive Function Timing Sequence. 254

E

ED Cell Circuit Schematic. 55
 Eight-Processor System Cabinet. 3
 Eight-Processor System Cabinet with Skins. 210
 Electrical Characteristics of a Transistor. 34
 Electronic Assemblies in the C-Pod—Front Left View. 246
 Electronic Components in the C-Pod—Front Right View. 249
 Exploded View of a CRAY-3 Module. 134
 Exploded View of Layers in a Module. 9

F

FE Cell Circuit Schematic. 56
 FF Cell Circuit Schematic. 58
 First Part of the AA Macro. 84
 FJ Cell Circuit Schematic. 59
 Flex Circuits Attached to Board Wirebonds. 177
 Flex Connector. 10
 Flex Connector Re-Flow Fixture. 176
 4.0 Micron Current Source—First Layer Metal. 35
 4.0 Micron Transistor—First Layer Metal. 33
 4.0 Micron Transistor—N+ Implants. 32
 4.0 Micron Transistor—Ohmic Metal. 32
 4.0 Micron Transistor—Semiconducting Channel. 31
 4.0 Micron Transistor—Recess Opening. 33
 Four-Processor System Cabinet and Control Pod. 2
 FQ Cell Circuit Schematic. 57
 Fractured Data File Name Conventions. 108

Fractured Stripes and Segments. 107
Front View of Eight-Processor System Cabinet. 209
Front View of a Four-Processor Tank. 208
Front View of Prototype One- or Two-Processor System Cabinet. 211

G

Generic IC Assembly. 80
Step One: The Generic IC Assembly is Created. 89
Gold Leads Bonded to a Portion of a Die. 131

H

Heavy Ion (N+) Implant. 119
Horizontal Power Buses for Logic Modules. 221
Horizontal Power Buses for Logic Modules from Module Side. 221
Horizontal Power Buses for Memory from Module Side. 220
Horizontal Power Buses for the Memory Modules. 220

I

Illustration of Stacking Layers. 185

K

The Kitting of Die for a Module. 173

L

Layers in a Module. 8
Linear Power Supplies in an Octant Assembly. 223
Logic Connectors Connected to a Corner Logic Board. 167
Logic Package Block Diagrams. 49
Logic Package Pinout. 51
Logic Package Power Grid and Cell Positions. 50

M

Macros and the Generic IC Assembly. 83
The Macro Lines for the Bonding Targets. 87
Maintenance Panel Controls. 244
Module Build Flow Chart. 172
Module Logic Connector with Flex Circuit. 169
Module Logic Plate. 152
Module Placement in the CRAY-3. 5

Module Power Blade. 162
Module Types Used in a CRAY-3 Octant. 6
Modules and Power Supplies in an Octant. 218
Modules Mounted in an Eight-Processor Cabinet. 216

N

Notcher Mechanism. 129

O

Ohmic and First Layer Metal in a 700 Ohm Resistor. 47
Ohmic Contact Metallization and Proton Isolation. 120
One Meg Memory Circuit Board. 149
Orientation of the FET Gates on a Wafer. 106

P

P Stack Logic Circuit Board. 144
Parts of a Macro Text Line. 82
Placement of the AC Macros. 86
Step Two: Placing the Required Logic Cells. 90
Plated Hole in a Logic Circuit Board. 142
Portion of a Logic Plate with Trace Lines. 153
Positioning of User Controls and Maintenance Panel. 241
Power Distribution in an Octant Assembly. 219
Power Distribution in Four Meg Memory Modules. 166
Power Distribution in Logic Modules. 164
Power Distribution in 256K Memory Modules. 165
Simplified Top View of a Vm Power Supply. 226
Simplified Top View of a Vn Power Supply. 227
Simplified Top View of a Vp Logic Linear Power Supply. 224
Simplified Top View of a Vp Memory Power Supply. 225
Pulse-Power Supply Block Diagram. 193
Pulse-Power Test Station. 191

R

Resistors on a Logic Module Resistor Plate. 157
Routed Portion of C Street West and Adjacent Avenue. 96
Step Three: Routing the Final IC Assembly. 94

S

SC Die Pinout. 71
Schottky Barrier and First Layer Metal Interconnect. 121
SD and SE Die Pinouts. 72

16-Up Mask Frame Layout. 103
Size Comparisons of the Memory Die. 70
Sizes of the Basic Components in a CRAY-3 Module. 11
Small Wires Soldered to Circuit Boards. 174
SM Die and Bypass Capacitors. 161
SM Die Pinouts. 73
Stacking a Module. 184
Stepped Die Patterns on a Wafer. 117
Stepping Sequence on a GaAs Wafer. 116
Sub-Floor Mounting of the C-Pod. 250
Sub-Floor Mounting of the System Cabinet. 212

T

Time Chart for Pulse-Power Testing. 194
Top Plate for a Four-Processor CRAY-3. 214
Top Plate for an Eight-Processor CRAY-3. 215
Top Plate for a Prototype One- or Two-Processor
CRAY-3. 217
Translator Box. 236
Twist Pin Defects. 182
Twist Pin Illustrations. 180
Twist Pin Seating. 187
Type A Diode (3 Square Micron)—N- and N+ Implants.
36
Type A Diode—Metal Layers and Recess Insulator. 37
Type C Diode/Capacitor—Ohmic Metal. 41
Type C Diode/Capacitor—First Layer Metal. 42
Type C Diode/Capacitor—Recess Insulator. 43

U

Understanding the Layer Title. 110

V

Vector Register Block Diagram. 67
Vp Power Distribution in a Module. 163
VR Read Timing. 69
VR Write Timing. 68

W

Wafer Dicing. 127
Wire Harness Connection Points to the Modules. 229
Simplified Diagram of Four-Processor Wire Harness. 230

Index

A

A Module 7
absolute coordinates 86
AC box 247
AC bulkhead 247
access time, memory 18, 70
acrylic lid 197, 208
address size, for IC masks 102
address units 107
AF cell 64
AG cell 58, 65
AH cell 66
AI cell 66
air bubbles 197
alignment marks 30, 102, 103, 109, 117, 118
alignment of FETs on wafer 104, 105
alignment pins 138
alignment process, Canon stepper 118
aluminum casting 210
and, logic 49, 52
anode 36, 38, 42, 43
Archive Header File 107
argon, use of 114
assembly pins 136
asymmetry, in clock amplifiers 62, 63
auto-leveling 243
Automatic Pin Insertion Machine 188–189
auto-router 153
avenue, in ICs 51, 79, 80

B

B Module 7
backgating 120
Background Processor 7, 17, 19–22, 190, 191, 196, 197, 218
ball formation 128
bandwidth, memory 17
bank, Common Memory 17, 18
bar code 103, 104, 108
bed of nails tester 141
beryllium-copper alloy, use of 168
beryllium-copper wire, use of 180
binary circuits 1
bits, correction 17
bits, data 17
board solder 174
bonding pads 13, 51, 60, 85, 86, 93, 104, 124, 128
bonding rate 129
boolean 79, 81, 88, 91, 97, 98, 124
buffer board 194
buffered outputs 57
buffers, data 17, 253
building a module 172
bulkheads 210
bus screws 163, 219
bypass capacitors 160

C

C Module 7
CA cell 60
CAD, see Computer Aided Design
Cadence 31, 74
cap, silicon nitride 118, 120
capacitive loading 62, 98
capacitor design 40–45
capacitor, electrical characteristics 45
capacitors, in pulse-power supplies 192
capacitors, reverse biased, in CRAY-3 ICs 50
capacitors, size of in CRAY-3 ICs 40–41
cathode 36, 38, 42
CB cell 60, 62
CC cell 63
cell library 48, 74–78
cell placement, in an IC 91, 92
cells 48, 81
cells, see also macros
channels, for cooling 231
chilled water 243, 248
chips 13
chips, see also integrated circuits
chromium 101
circulation pump 231, 237
clamping function, of diodes 36
clamping, voltage 53, 54, 58, 66
clean room 101, 139
clear, data on mask layers 109
clock amplifier design 60–64
clock inputs 51
clock phase time 4
clock skew 62
clock speed 194
clock, invert 4, 60, 95
clock, normal 4, 95
clock, odd phases 4
Common Memory 7, 17–18, 197, 228
complementary, logic output 54, 55
Composite, Board Stack 146
Computer Aided Design (CAD) 31, 74, 81, 82
conditioning filters 231
connector, 2 by 22 9, 145
connector, flex 9, 167–169, 175, 218, 229
contact immersion cooling 231
contact probes, use of 191
Control Pod 2, 197, 231, 237–250
controllers, peripheral equipment 17, 251, 252
cooling 231–234
cooling channels 231
copper, use of 139, 141, 150, 158
CPU 7
CPU, see also Background Processor
Cray Research, Inc. 251
CRAY-2 1, 15–16, 231
CRAY-2, use of in design of CRAY-3 99, 153

CrayDraw 31, 74, 139, 145, 153
crystalline structure of gallium arsenide 104, 113
current rating, of pulse-power supplies 192
current source design 35, 52, 53, 55
cutter, for jumpers 188
cycle time, memory 18, 70, 71

D

D cell design 52–54
D Module 7
data rotation 104
data transfers 256, 257
data, on glass masks 102–104
DC bulkhead 245, 248
Dead Start 194, 237
defects in twist pins 181
delay chains, long 64
density, logic gate 27
depletion mode 113
design rule checks (DRCs) 98, 99
diagnostic software 195, 196, 199, 202
dicing 126–127
die 13, 48, 104, 127, 132, 147, 173, 177, 178, 202, 203
Die Attach 177
die stamper 179
die, see also integrated circuits
dielectric 33, 114, 121, 139
differential coupling 60
dimensions, bypass capacitors 160
dimensions, C-Pod 237
dimensions, drilled hole 155, 158
dimensions, drilled holes 150
dimensions, IC packages 83
dimensions, logic plate traces 150
dimensions, logic signal trace line 12
dimensions, module 7, 133
dimensions, of circuit boards 139
dimensions, of drilled hole 140
dimensions, of GaAs wafer 113
dimensions, of gold leads 137
dimensions, of plated holes 141
dimensions, of printed circuit boards 141
dimensions, of trace metal 142
dimensions, of Ultem pieces 135
dimensions, power and resistor plates 185
dimensions, printed circuit boards 11
dimensions, reticles 116
dimensions, System Cabinet 1, 207
dimensions, trace lines in IC 94
dimensions, vias in ICs 96
dimensions, wafers 125
diode design 36–40
diode, type A 36, 39
diode, type B 38, 39

Disk Array 257
 Disk Subsystem 251
 DITMCO tester 141
 double implant 32
 drain current 34, 35
 drain element 31, 33
 drain voltage 34, 35
 drawing layer 87
 duty cycle 193

E

E cell design 54–55
 E Module 7
 E-beam, see also electron beam 101, 102
 Electroglas 124
 electroless copper process 140
 electron beam evaporation 121, 122
 electronic module, of C-Pod 245
 emitter coupled logic (ECL) 48, 235
 epoxy glue, use of 168, 175
 epoxy, use of 139
 equipment used in full-power testing 196
 equipment used in pulse-power testing 192
 errors, in testing of modules 195
 ESD protection 62, 64, 65
 Extract 98

F

F cell design 56–57
 F Module 7
 failure analysis 203
 failure analysis codes 204
 false, logic state 229
 FAST testing 57, 59, 124–125
 FF and FG cell design 58
 field effect transistor (FET) 29, 54, 104, 121
 field numbering 112
 field, dark 102
 field, light 102
 FIFO 258
 filters 237
 first layer metal 33, 35, 36, 50, 96, 121
 FJ and FK cell design 59
 flat of wafer 105, 106, 110, 113
 flex connector, advantages of 168
 flow blockers 183, 231
 flow meters 243, 248
 Fluorinert 4, 137, 154, 237, 243, 244, 248
 foreground channel loop 252, 258
 Foreground Processor 7, 17, 23–25, 190, 191, 193, 195, 196, 197, 218, 252, 258
 forward biased, of capacitor 45
 forward current, of diode 39, 40
 forward resistance, of diode 39
 fracturing 104–108

front panel user controls 241
 full-power testing 196–202
 full-power testing, advantages of 196

G

G Module 7
 gallium arsenide 2, 13, 15, 29
 gate array based 29
 gate current 29
 gate element 31
 gate metal 33, 35
 gate voltage 29, 34
 gate, logic, density 27
 gate, orientation on wafer 105
 gates, logic, type of used in CRAY-3 48
 glass plates, for ICs, see also masks 101, 184
 gold germanium, use of 115
 gold, use of 114, 115, 121, 122, 128, 141, 150, 158, 168, 180
 grid, 0.1 micron 31, 83
 guide pins 186

H

H Module 7
 hand placement of cells in ICs 97
 hand routing of ICs 97
 heat 3
 heat exchanger 231, 237
 high-power microscope, use of 203
 high-speed cards 235
 HIPPI 235, 257, 258
 horizontal power buses 220, 221
 hoses 237
 Hughes Automatic Pin Insertion Machine 188
 Hughes bonding machine 128, 179
 Hughes Semi-Automatic Pin Insertion Machine 186
 hydrogen, use of 120

I

I/O Module 190, 196, 197, 218
 IC assembly 79–89
 IC Master 99, 104
 IC, test 104
 Idss 34
 impedance 142, 150
 inductive probes, use of 191
 infrared microscope, use of 195
 ingot 113, 120
 inputs, maximum number for IC package 81
 inscriptions, IC names 112
 integrated circuits 13, 79, 101, 113, 135, 136
 inverting amplifier 49, 54, 60
 ion implantation 114

J
jumpers 136, 140, 180
jumpers, defects 182
jumpers, inspection of 181
jumpers, mechanical 144
jumpers, number per module 186
jumpers, power 162

K
kitting 173

L
L Module 7
lamination of circuit boards 140
laser welding 181
latch failure 64
layer title 110
layers in a module 8, 9
layers, in circuit board 143
layers, of memory circuit boards 147
leader-less twist pin 188
leakage, in integrated circuit design 33, 120
level sensors 243, 244
level shifting 36, 54, 57, 58, 65, 66
lift-off process 121, 122
linear accelerator 114
loading, internal 49, 54, 55
logic gates 36
logic package design 48-64
logic plates 8, 12, 134, 150-153, 171, 183
low-speed cards 235

M
M Module 7
machine clock period 4
Macintosh 99, 125, 193, 194, 195, 197
macros, nesting of 82
macros, see also cells 48, 81, 82, 83
maintenance panel 244
masks, circuit boards 139
masks, integrated circuits 30, 81, 101-112, 115
maximum field strength, of resistor 46
maximum number of inputs for IC 81
maximum number of outputs for IC 81
MEBES 102, 104, 107
memory, common 70
memory, local 70
Metal Oxide Semiconductors (MOSFETS) 29, 192
Metal Semiconductors (MESFETS) 29, 113
metallization variation 29
module 7, 133, 171
module build 172
module clock 254
module layers 8

module pallet 186
module repair 202-203
module stacking 183-186
module testing 190-202
module, power distribution in 163-166
motor generator (MG) 222, 244
motor generator exciter 243
motor starters 247

N
N- implant 31, 32, 42, 119
N+ implant 31, 42, 46, 66, 119
National Center for Atmospheric Research (NCAR) 3
Neodymium Glass laser 181
nickel, use of 141, 150, 158
nitrogen, use of 114
noise, high frequency 160
non-metallic bolts 197, 208
normal, logic output 52, 54, 55

O
octant 4, 210, 213
octant valves 243
octant, use of term 5
Ohmega layer 134
ohmic contact 36, 39, 46, 115
ohmic metal 32, 33, 127
one, logic state 48, 194, 229
operating temperature 231
or, logic 49, 53, 54, 56
orientation of FET on GaAs wafer 114
O-ring 208
oscillator 3
outputs, maximum number for IC package 81

P
pad assignments 93
passivation layer 104, 109, 123
passivation openings 30
path resistance, diodes 36
PECVD 114, 121
Perkin-Elmer E-Beam 101
phase shift, in clock circuits 60
photo-optical sensor 189
photoresist 114, 115, 121, 139, 140
photoresist, retro-grade 122
pin insertion rate 188
pin insertion, rate of 189
pinch-off voltage 34, 35
pinning 186-189
pinout, memory die 71-73
piping 237
PlaceCells 81
plasma enhanced chemical vapor deposition 114

plasma etch 121
 plating uniformity 158
 platinum, use of 115
 polarities, logic 49
 ports, memory access 17
 position of ICs on wafer 110
 power blades 9, 133, 158, 162–163, 183, 185, 186, 219
 power buses 162, 219
 power distribution 163, 219
 power grid, in ICs 50, 51
 power pins 138, 185
 power plates 8, 12, 134, 158–159, 171, 183
 power plates, layers 158
 power supplies 1, 207, 222–227
 power supplies, for pulse-power testing 192
 power surge protection 64
 power, watts 3, 4, 48, 231
 pressure transducers 243
 printed circuit board manufacturing 139–141
 probe pads 104
 Process Control Monitor (PCM) 103, 104, 127
 program code 17
 propagation delay, see also speed 54, 55, 62
 protection circuitry 192
 protective overlay 135
 proton isolation 120
 puller, for insertion of jumpers 188
 puller-cutter cycle 187
 pulse-power testing 190–195
 pulse-power testing, advantages of 190

R

RAID 251, 257
 RC time constant, in clock amplifiers 60
 read timing 69
 recess insulator 43
 recess opening 33
 referenced coordinates 85
 reservoir 233, 237
 resistivity, co-planar trace on logic plate 150
 resistivity, of GaAs 114
 resistivity, of implant 119
 resistor design 46–47
 resistor plate 8, 12, 56, 154–157, 171, 183, 185
 resistor plate, layers 154
 reticles, see also masks 113
 reverse biased capacitance, of diodes 40
 reversed biased mode, of capacitor 45
 Route 81, 98
 routing, of ICs 88–97
 RS-232 248
 RS-485 248
 RTA 120

S
 saturation, in design of resistor 46
 SC die 70
 Schottky barrier, see also Schottky diode 115
 Schottky diode 29, 32, 36, 113
 Schottky metal, see first layer metal
 Schottky-Diode FET Logic, see SDFL
 SD die 70
 SDFL 48
 SECEDED 18
 second layer metal 50, 51, 122
 seed layer 140
 segments, fractured 106, 107
 semiconducting channel 31
 semiconducting channel, length 31
 series resistance, in capacitor design 44
 SF die 70
 sheet resistance 32, 46
 shims 231
 shrinkage factor 31, 46
 sidegating 120
 signal generator 247
 signal termination 66
 silane, use of 115
 silicon 2, 15, 114, 119
 silicon dioxide, use of 115, 123
 silicon nitride, use of 114, 118
 simplex channel 258
 SM die 71
 solder quick tape, use of 175
 solder sleeve 230
 soldering 174
 soldering, Pulsed Reflow 175
 source element 31, 33
 source follower 52, 54, 55
 spacers 137, 175, 183
 speed 48, 124
 speed, of transistor 33
 spot size 102
 sputtering 114, 115
 square, of diode 36
 SRAM 18, 65, 66
 Stacking Fixture, 16-Up 183, 184
 stacks 17, 135, 137, 138
 stacks, control 17
 stamper, for die 179
 stamping force 179, 203
 standard cell based 29
 stepper, Canon 101, 103, 104, 116–118
 strain reliefs 175
 streets, in ICs 51, 79, 80
 stripes, fractured 106, 107
 sub-floor 208
 substrate 29, 33, 113, 115
 symmetry, in clock amplifiers 62
 System Cabinet 1, 2, 4, 207–217

1026 21207

system clock 247, 252
system reset 243
System Test Interface

T

tank 6, 207, 208, 210, 213
tank, see also System Cabinet
Teflon, use of 167
temperature, operating 231
terminating resistor 56, 58, 134
terminating resistors, number on resistor plate 154
test station 190
testing, of board artwork 146
testing, of IC artwork 98
thermocouples 243
thermosonic ball bonding 128
threshold voltage 48
titanium, use of 115, 117, 118, 122
Toshiba 160
touch screen display 241
trace lines, in ICs 51, 96
transistor, design 31–35
transistor, electrical characteristics 34
Transistor-to-Transistor Logic (TTL) 194
translator cards 235
transmission line 48, 56
true, logic state 229
twist pins, see jumpers

U

Ultem overlay 135, 183
ultrasonic welding 128
ultraviolet light 116, 118, 139
Universal Stacking Fixture 183
UNIX 237, 247
up/down pump 237

V

Vector Register 13, 50, 60, 67–69, 74
Vg 50, 51, 64
Vgs 34
vias 96, 121, 140, 146
Vm 192, 222, 244
Vn 50, 51, 192, 222, 223, 244
Vn power bus, in ICs 93
Voh 57, 58, 59
Vol 57, 59
voltage droops 160
voltage regulation 192
Vp 50, 51, 64, 192, 222, 223, 244
VPlat 98

W

waffle pack block 178
waffle packs 125, 127, 128, 173

wave length of light 118
wire harness 218, 227, 228–230
wirebonds 145, 152, 175
wires, twisted pair 167
word 17, 251
write timing 68

Z

zero, logic state 48, 194, 229