CRAY SV1™ Series and CRAY J90se™ Series Programming Differences

(CRAY SV1 Series)

108-0226-001

Silicon Graphics Private

SiliconGraphics
Computer Systems

Any shipment to a country outside of the United States requires a letter of assurance from Silicon Graphics, Inc.

Requests for copies of service publications should be directed to:

SILICON GRAPHICS, INC.
Customer and Professional Services - Logistics
1100 Lowater Road
P. O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Comments about this publication should be directed to:

SILICON GRAPHICS, INC.
Customer and Professional Services
Service Publications and Training
890 Industrial Blvd.
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

CRAY SV1 Series and CRAY J90se Series Programming Differences

108-0226-001
CRAY SV1™ Series
Last Modified: December 1998

Figures

Tables

Record of Revision

December 1998

Original printing.

CRAY SV1 System Description

The CRAY SV1 series of computers (refer to Figure 1) offers a low-cost, air-cooled, binary compatible extension of the popular CRAY J90se series computer systems. The CRAY SV1 hardware architecture permits a scalable system configuration from 4 CPUs to 32 CPUs in 4-CPU increments with comparable memory sizes from 2 Gbytes to 32 Gbytes. Not all combinations are marketed.

This document describes the programming differences between the CRAY J90se series computer systems and the CRAY SV1 series computer systems.

Note: Refer to the *System Programmer Reference* manual, publication number CSM-0301-0B0, for comprehensive programming information for the CRAY J90 and CRAY J90se mainframes.

Figure 1. CRAY SV1 Series Systems

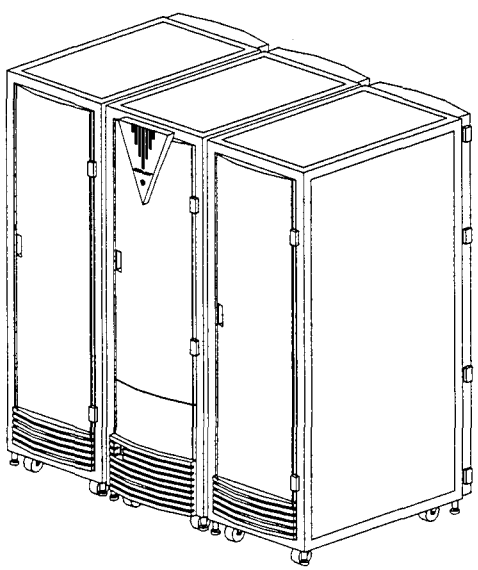

CRAY SV1-1A System

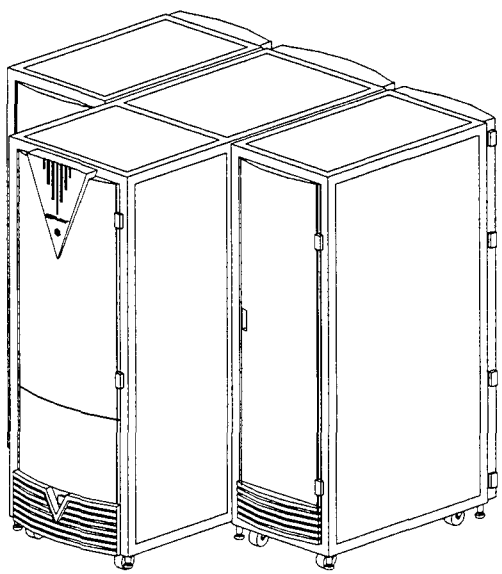

CRAY SV1-1 System

CRAY SV1 System Characteristics

The CRAY SV1 series systems have the following characteristics relative to the CRAY J90se series systems:

- The PC cache is replaced by a larger external cache, which is shared by both scalar and vector memory requests.
- The dual vector pipes replace the single vector pipe, but they retain the same maximum vector length of 64 elements.
- The PV array combines the CRAY J90 series scalar PC array and the vector VU array in a single, dual-pipe CPU array.
- The processor module clock frequency is 300 MHz instead of 100 MHz.
- The CPU interfaces with the VA and VB arrays via the CA cache arrays at 300 MHz instead of 100 MHz. Each CPU has two CA arrays.
- The cache size for each CA array is 16 Kwords (32 Kwords per CPU) instead of a single 128-word cache for each CPU.
- The scalar and vector functional units share the floating-point functional units of one pipe instead of using separate functional units. A scalar floating-point request is granted immediately even if the functional unit is busy. A vector operation in that functional unit is delayed by 1 clock period (CP) while the scalar instruction executes.
- The CRAY SV1 system uses the GigaRing I/O architecture. Customers have the option to use their existing VME I/O systems, but VME I/O systems are not configured with new CRAY SV1 systems.

CRAY SV1 Instructions

- The test and set instruction (0034*jk*) invalidates the CRAY SV1 CPU cache to enhance Autotasking performance. The invalidation of cache by the 0034*jk* instruction is disabled when the "disable cache invalidate" bit is set in the exchange package.

- A scalar read request to memory results in an 8-word request to memory for the 8 words associated with the same base address. The other 7 words are put into cache but they are not delivered to the CPU.

- The scalar bit matrix multiply (BMM) execute instruction (070*ij*6) is implemented to be issued as a vector instruction and to execute in the vector unit. The execution of this instruction is enabled by setting a configuration bit. This instruction is normally disabled. If the 070*ij*6 instruction issues when the configuration bit is not set (disabled), then the instruction executes as a floating-point reciprocal instruction.

- The vector leading zero instruction (174*ij*3) is implemented to share the floating-point reciprocal/pop/parity functional unit. The execution of this instruction is enabled by setting a configuration bit. This instruction is normally disabled. If the 174*ij*3 instruction issues when the configuration bit is not set (disabled), then the instruction executes as a floating-point reciprocal instruction.

- The vector load BMM register instruction (1740*j*4) is implemented in hardware to load up to 64 BMM registers. The BMM functional unit shares ports with the floating-point add functional unit. The execution of this instruction is enabled by setting a configuration bit. This instruction is normally disabled. If the 1740*j*4 instruction issues when the configuration bit is not set (disabled), then the instruction executes as a floating-point reciprocal instruction.

- The vector BMM execute instruction (174*ij*6) shares ports with the floating-point add functional unit. The execution of this instruction is enabled by setting a configuration bit. This instruction is normally disabled. If the 174*ij*6 instruction issues when the configuration bit is not set (disabled), then the instruction executes as a parity instruction.

CRAY SV1 Processor Module

Each CRAY SV1 processor module includes 4 CPUs. Each CPU consists of one PV application-specific integrated circuit (ASIC) array and two CA arrays. The PV/CA interface is organized on a section basis in which two sections use common buses between the arrays. The interface between the PV array and each CA array consists of two address and two request buses, one write data bus, and two read data buses.

The following paragraphs describe the section assignment for these buses:

PV and CA0: Sections 0 and 2 share one bus, and sections 5 and 7 share the other bus. The write data bus is common to all four sections.

PV and CA1: Sections 4 and 6 share one bus, and sections 1 and 3 share the other bus. The write data bus is common to all four sections.

This bus sharing by sections is designed for best stride-of-2 and stride-of-4 performance. A cache reference within the CA takes 2 cycles. Thus, two different section references share the same bus.

The peak vector performance of the CRAY SV1 processor is enhanced to 1,000 MFLOPS. The CRAY SV1 processor module is a field upgradable replacement for a CRAY J90 series or a CRAY J90se series processor module.

The current CRAY J90se series processor interfaces with the CRAY SV1 processor using a 10-ns CP. In addition, the CRAY J90se series processor CI and JS arrays are used on the CRAY SV1 processor with the same 10-ns CP interface to the faster arrays. The current CRAY J90se series MC array is also used.

ASIC Differences

Table 1 shows the timing differences between the various ASICs in a CRAY SV1 processor and the ASICs in a CRAY J90se processor.

Figure 2 provides block diagrams of the ASIC configurations.

Table 1. ASIC Differences

CRAY SV1 System			CRAY J90se System		
ASIC	Frequency	CP	ASIC	Frequency	CP
Memory	100 MHz	10 ns	Memory	100 MHz	10 ns
to VB	300 MHz	3.3 ns	to VB	100 MHz	10 ns
to CA 32-Kword cache	300 MHz	3.3 ns	to PC+ 128-word cache	200 MHz	5 ns
to PV	300 MHz	3.3 ns	to VU	100 MHz	10 ns
to CA	300 MHz	3.3 ns	to VA	100 MHz	10 ns
to VA	300 MHz	3.3 ns	to Memory	100 MHz	10 ns
to Memory	100 MHz	10 ns			

Figure 2. CRAY SV1 and CRAY J90se ASIC Configurations

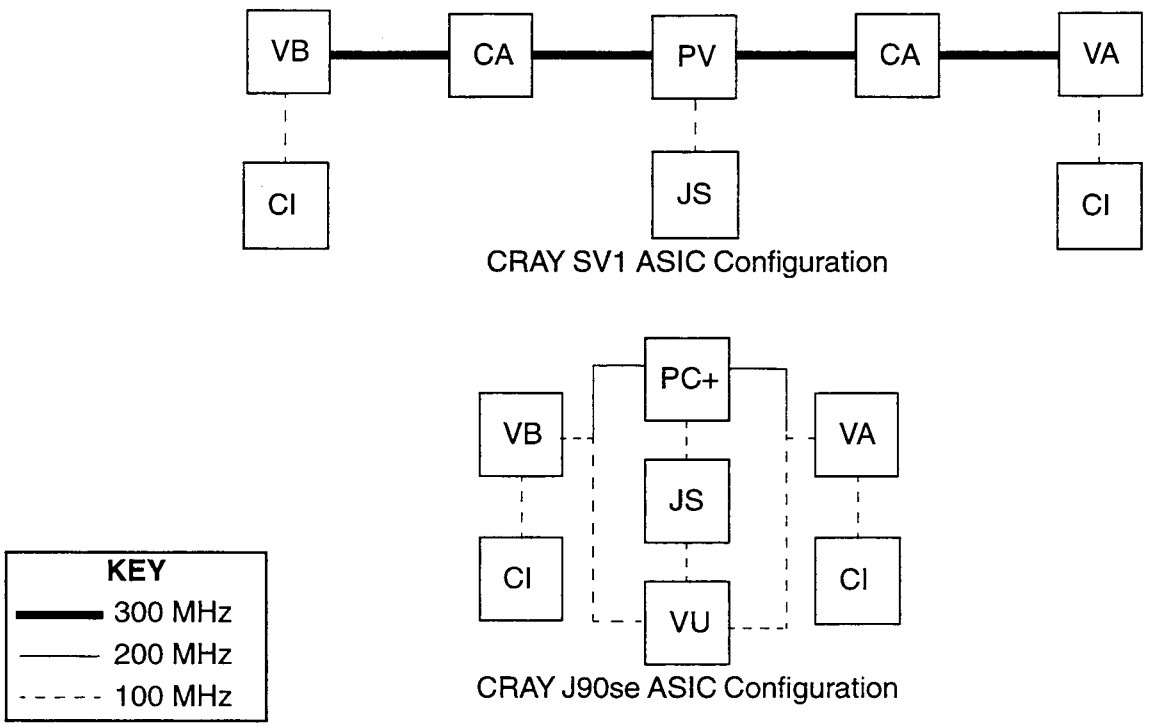

The new CRAY SV1 arrays are as follows:

- PV: The combined scalar and vector dual-pipe CPU array
- CA: The cache array with 16 Kwords per CA
- VA (VAB): The CPU memory interface for read/write requests and for write data to memory
- VB (VAB): The CPU-to-memory interface for read data from memory

  The VA and VB functions reside in the same VAB array type. Eight VAB arrays are required for a processor module: four function as VA arrays and the other four function as VB arrays. The CRAY SV1 VA and VB arrays interface with memory through the backplane. The VAB array interfaces with the backplane at 100 MHz, and it interfaces with the CA cache array at 300 MHz.

Cache Line Size

The line size of the cache is 1 word per line. For scalar references, the cache operates as if the line size were 8 words (refer to Figure 3).

The time for cache hits is 9 CPs on the cache chip. The total hit time is approximately 25 CPs; the miss time, or memory reference time, is approximately 90 CPs.

Figure 3. Cache Data Paths

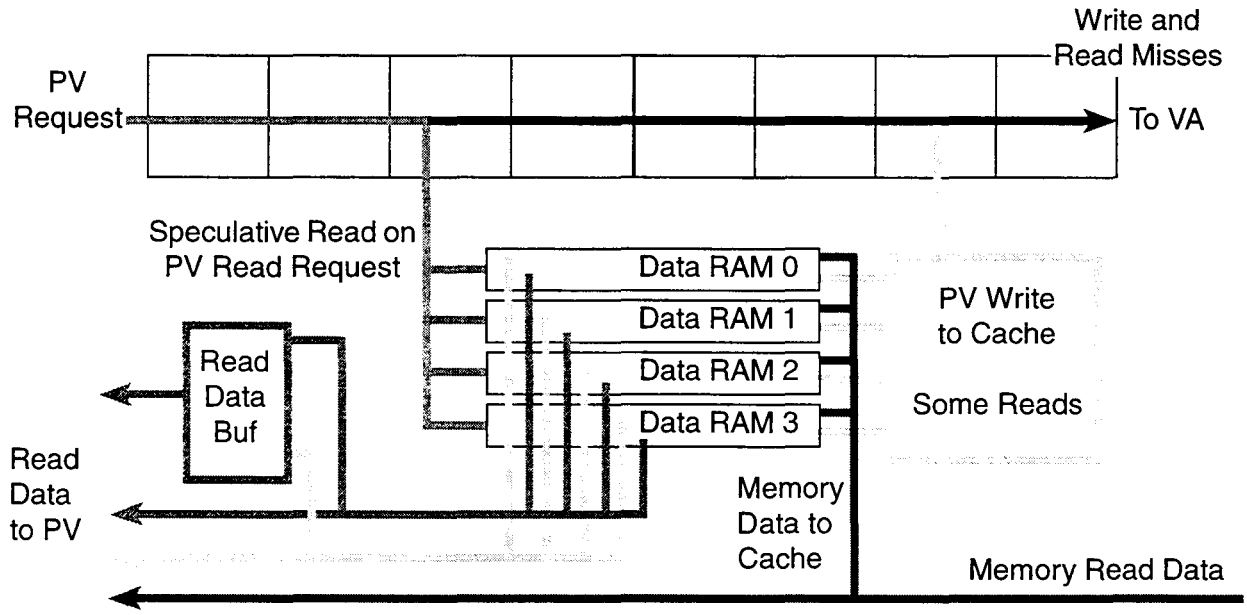

Programming Differences

CRAY SV1 series computer systems are binary compatible with the CRAY J90se series computer systems for user applications.

The CRAY J90se series design includes scalar first-level cache and a 0016*j*1 instruction (IVC). This instruction enables any CPU in the system to invalidate the cache of another CPU in the 32-CPU system.

The CRAY SV1 design includes the following changes and additions:

- The CRAY J90se series 128-word scalar cache is removed and replaced by a 32-Kword cache per CPU. This cache is located on two CA arrays and is used for all instruction, scalar, B, T, and vector data. This cache is invalidated by the 0016*j*1 instruction (same as CRAY J90se series), by the 0034*jk* (SM*jk* 1,TS) test and set instruction, and before and after every exchange operation.

- The 0034*jk* (SM*jk* 1,TS) test and set instruction has an added feature in the CRAY SV1 system: the issue of this instruction invalidates cache for that CPU (when the Enable Cache Invalidate bit is set in the exchange package).

- The CRAY SV1 system is a dual-pipe vector design that retains the vector length limit of 64. This dual-port design is similar to the CRAY J90se series, but the ports are expanded from "A and B" to "A and A´ and B and B´" to match the dual-pipe design. The ports are internal to the PV array (CPU chip) and do not have data error correction directly associated with them. The CPU data error correction occurs on the cache (CA) arrays at their interface with memory (and is not directly associated with ports). The port A and port B maintenance instructions, 001501 and 001511 respectively, will execute identically. Either one will disable CPU error correction on the CA arrays. The 001521 instruction will still disable I/O error correction in the CI array, and the 001541 and 001551 instructions remain unchanged.

- The CRAY SV1 system has both the scalar design and the dual-pipe vector design located in the same CMOS array (chip). To save space within this array, the scalar floating-point instructions execute in the vector pipe 0 floating-point units. These scalar floating-point instructions are issued and executed without regard to any current vector instruction activity. (The scalar floating-point unit uses a clock cycle from the vector floating-point unit.)

- The CRAY SV1 design includes a dual-pipe bit matrix multiply (BMM) unit that is compatible with previous Cray system architectures. The CRAY SV1 design includes the following additional instructions:

Table 2. CRAY SV1 BMM Instructions

Instruction	Description
002210 (CBL)	Clear Bit Matrix Loaded bit (BML) in the status register and exchange package. Note that the BML bit replaces the PS bit just as in previous Cray system architectures.
070*ij*6 (S*i* S*j**BT)	Transmit the bit-matrix product of (S*j*) and the matrix BT to S*i*.
1740*j*4 (BMM LVj)	Transmit V*j* elements to BT matrix. Note that this load of the 64-register BT matrix clears out all locations of the matrix that are not loaded. To save space in the array, the load of the BT Matrix requires 8, 16, 24, or 32 cycles, depending on vector length (instead of just VL/2 cycles).
174*ij*6 (V*i* V*j**BT)	Transmit bit matrix product of (V*j*) and (BT) to V*i*.

- The CRAY SV1 system, just like CRAY J90se series, is designed with hardware configuration bits that define and control the system hardware. A new configuration bit enables the BMM instructions in the CPU. When BMM is not enabled, the 002210 instruction is a no-op and the other three instructions operate the same way as in the CRAY J90se series systems. (The configuration bits that define the system are loaded from a file in the same manner that the CRAY J90se series systems use.)

Vector Leading Zero Instruction

The CRAY SV1 design also includes a new instruction, 174*ij*3 (V*i* ZV*j*), which is the vector leading zero instruction. This dual-pipe operation is executed in the reciprocal/pop/parity/leading zero functional unit.

Exchange Package

Figure 4 shows the exchange package format for the CRAY J90se series and CRAY SV1 series systems.

Two exchange package bits designate the processor type: a CRAY J90 processor, a CRAY J90se processor, or a CRAY SV1 processor.

Exchange bits in word 7 (bits 30 and 31) are saved in bit positions 62 and 63 of the full word that is stored to memory. In the CRAY J90 series systems, both bits are 0's. In CRAY J90se series systems, bit 63 is 0 and bit 62 is 1.

The UNICOS operating system compensates for the differences in performance between CRAY J90 processors, CRAY J90se processors, and CRAY SV1 processors. Also, some library routines have been altered to accommodate the differences in processor performance.

Table 3 describes the bits and lists the word and bit positions. The CRAY SV1 exchange package defines additional bits as follows:

Table 3. CRAY SV1 Exchange Package Differences

Word and Bit	Description
Hardware numbering (refer to Figure 4)	
Word 6 bit 42	The Program State (PS) is redefined as the Bit Matrix Loaded (BML) bit when BMM is enabled for that system. This is consistent with previous Cray architectures.
Word 7 bit 32	CRAY J90se series Cache Enable (CE) bit redefined as CRAY SV1 Enable Cache for Data (ECD)
Word 7 bit 33	CRAY SV1 Enable Cache for Fetch (ECF)
Word 7 bit 34	CRAY SV1 Enable Cache Invalidate (ECI). Allow 0034*jk* to invalidate Cache.
Word 7 bit 35	CRAY SV1 Disable Scalar Cache Hit counts to PM
Word 7 bit 36	CRAY SV1 Disable Fetch Cache Hit counts to PM
Word 7 bit 37	CRAY SV1 Disable B/T Cache Hit counts to PM
Word 7 bit 38	CRAY SV1 Disable Vector Cache Hit counts to PM
Word 7 bits 63:62	These two hardware bits, which are defined by the PV array, are written into memory to identify the type of system. [63:62] = 2b00 for a CRAY J90 series system [63:62] = 2b01 for a CRAY J90se system [63:62] = 2b10 for a CRAY SV1 system

Figure 4. CRAY J90, CRAY J90se, and CRAY SV1 Series Exchange Package

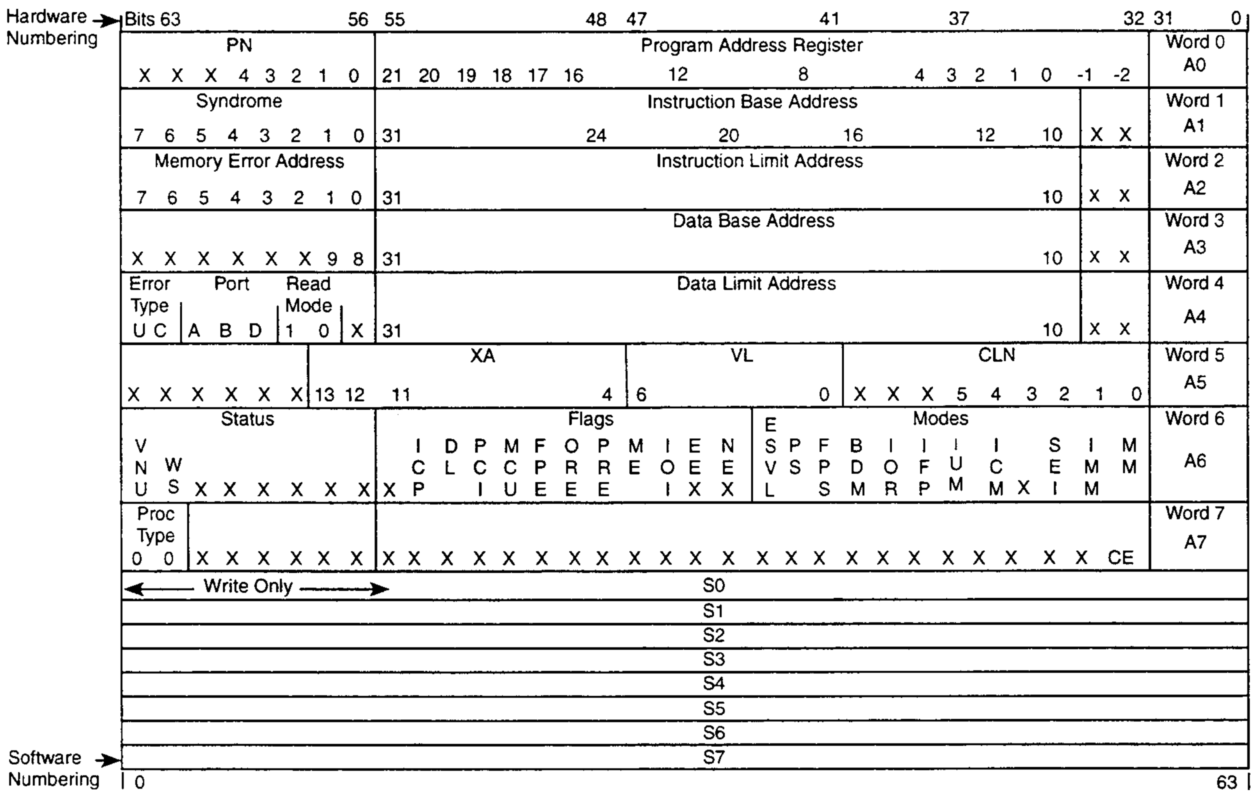

CRAY J90 and CRAY J90se Read Mode Error Reporting

Bit 1	Bit 0	Port A	Port B	Port D Channel
0	0	EX	Fetch A	n+1
0	1	B	T	n+3
1	0	Vector	Vector	n+5
1	1	A/S	Fetch B	n+7

n = Processor board number X 10 + 20

Processor Type

Bit 31	Bit 30	Type
0	0	CRAY J90
0	1	CRAY J90se
1	0	CRAY SV1
1	1	Unused

CRAY SV1 Read Mode Error Reporting

Bit 1	Bit 0	Port A	Port B	Port D Channel	Operation with Error
0	0	1	0	0	Exchange
0	0	0	1	0	Fetch
0	1	1	0	0	B
0	1	0	1	0	T
1	0	0	0	0	Vector
1	1	1	0	0	A or S
0	0	0	0	1	Channel n+1
0	1	0	0	1	Channel n+3
1	0	0	0	1	Channel n+5
1	1	0	0	1	Channel n+7

n = Processor board number X 10 + 20

Notes:

- CRAY SV1 exchange package words 0 through 5 are identical to CRAY J90se series systems.
- CRAY SV1 exchange package word 6 is similar except that Mode bit PS (Program State) is replaced by BML (Bit Matrix Loaded) in bit 42 as shown below.
- CRAY SV1 exchange package word 7 has additional bits defined as shown below and defined in Table 3.

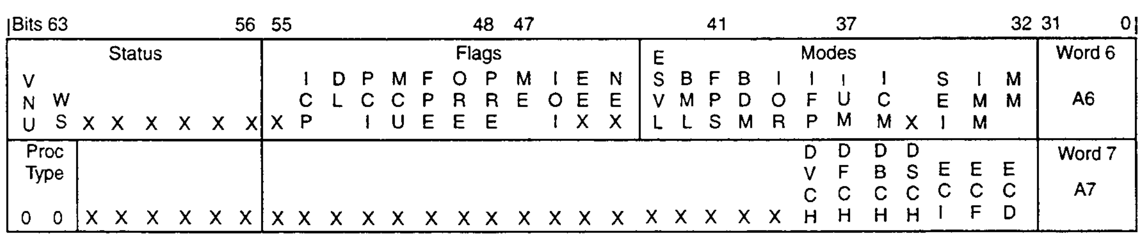

Performance Monitor

The CRAY SV1 performance monitor includes a change in the definition of a *cache hit*. Bits 35 through 38 of exchange package word 7 control the types of cache hits that are counted. All types of cache hits are counted when these 4 bits are clear.

Error Detection and Correction

The CRAY SV1 system includes a change that is related to error detection and correction. A word from memory that contains a double-bit error will be returned as all zeroes (unless error correction is disabled).

Vector Functional Units

The following CRAY SV1 vector functional units are grouped together to share data buses from the vector registers:

Floating-point Add unit and	Bit Matrix Multiply unit
Floating-point Multiply unit and	Second Vector Logical unit (VL2)
Floating-point Reciprocal unit and	Pop/Parity/Vector Leading Zero
Vector Add unit	None
Vector Logical unit	None
Vector Shift unit	None
Vector Store Data unit	None (The 076 instruction takes a cycle from these instructions)
Vector Index Data unit (for Scatter/Gather)	None (The 076 instruction takes a cycle from these instructions)

The CRAY SV1 design uses the same floating-point units as the CRAY J90se series systems. Functional unit latency depends on the depth of the functional unit. Table 4 lists the functional unit latency for each unit:

Table 4. Functional Unit Latency

Function Unit	Data Register Depth	CRAY J90se and CRAY SV1 Differences
Vector Floating-point Add	6 registers deep	Same
Vector Floating-point Multiply	7 registers deep	Same
Vector Floating-point Reciprocal	14 registers deep	Same
Scalar Floating-point Add	6 registers deep	+ 2 cycles
Scalar Floating-point Multiply	7 registers deep	+ 2 cycles
Scalar Floating-point Reciprocal	14 registers deep	+ 2 cycles
Bit Matrix Multiply	2 registers deep	New
Vector Logical 2	2 registers deep	+ 1 cycle
Vector Pop/Parity	3 registers deep	Same
Vector Leading Zero	2 registers deep	New
Vector Add	2 registers deep	Same
Vector Logical	2 registers deep for logical and 3 registers deep for compress Index and 2 registers deep for Vector Mask	+ 1 cycle - 2 cycle - 2 cycle
Vector Shift	4 registers deep for Left Double and 3 registers deep for all other	+ 1 cycle + 1 cycle

The CRAY SV1 scalar floating-point instruction results are available 2 cycles later than in the CRAY J90 series systems. The vector instruction results, from the above functional units, are available for use by the next vector instruction after 2 cycles (+1 cycle more than the CRAY J90se series systems).

All CRAY SV1 vector functional units are pipelined. Generally, one dead cycle exists between the last data sent to a functional unit and the first data for a following instruction to be received by the same functional unit. The same is true for a companion functional unit that shares the same data bus from (not to) the vector registers. This is also true for floating-point multiply to VL2 and floating-point reciprocal to pop/parity/leading zero transitions.

The BMM load is the exception. Generally, one dead cycle exists for vector register usage for source-to-source register reservations or result-to-result register reservations between instructions. No instruction issue check occurs for source-to-result or result-to-source register usage (because of chaining and tailgating).

Cache Latency Timing

Cache is physically located on two CA arrays per CPU. The latency of these cache arrays (only) is 9 cycles, effectively. The current best estimates for latency for some of the operations with cache are listed below:

- Scalar latency timing with cache from CIP instruction issue:

AS data from cache	Approximately 22 cycles
Fetch from cache to CIP	Approximately 28 cycles

- Vector load latency timing with cache from CIP instruction issue:

VR data from cache with no chaining/tailgating	Approximately 27 cycles
or with chaining/tailgating[1]	Approximately 32 cycles

[1] There is increased latency associated with Chaining and Tailgating above, but the overlap in operations due to Chaining and Tailgating may permit much earlier issue of the instruction by the Vector Issue Register (VIR) than would otherwise occur.

- Vector gather latency timing with cache from CIP instruction issue:

VR data from cache with no chaining/tailgating	Approximately 30 cycles
or with chaining/tailgating[1]	Approximately 34 cycles

[1] There is increased latency associated with Chaining and Tailgating above, but the overlap in operations due to Chaining and Tailgating may permit much earlier issue of the instruction by the Vector Issue Register (VIR) than would otherwise occur.

- The current best estimates for latency for some of the operations with memory are listed below. The latency is listed for a 300-MHz system clock.

- Scalar Load latency timing with memory from CIP instruction issue:

A/S data from Memory at 300 MHz	Approximately 102–105 cycles
Fetch from Memory to CIP at 300 MHz	Approximately 108–111 cycles

- Vector Load latency timing with memory from CIP instruction issue:

VR data from Memory with no chaining or tailgating at 300 MHz	Approximately 107–110 cycles
or with no chaining or tailgating[1] at 300 MHz	Approximately 112–115 cycles

[1] There is increased latency associated with Chaining and Tailgating above, but the overlap in operations due to Chaining and Tailgating may permit much earlier issue of the instruction by the Vector Issue Register (VIR) than would otherwise occur.

- Vector Gather latency timing with Memory from CIP instruction issue:

VR data from Memory with no chaining or tailgating at 300 MHz	Approximately 110–113 cycles
or with no chaining or tailgating at 300 MHz	Approximately 115–118 cycles

- Exchange operation latency timing to first CIP instruction issue:

Exchange data from memory at 300 MHz	Approximately 217–223 cycles

The following assumptions were used to determine memory latency:

- CPU clock frequency of 300 MHz
- Memory module clock frequency of 100 MHz
- Synchronization at the 300 MHz to 100 MHz interface occurs over a 3-CPU-cycle window (1, 2, or 3 cycles required), so latency can vary accordingly for each request to memory.
- Memory latency depends on the ratio between the CPU and memory clock frequencies. Although the memory latency in memory clock periods is a constant (approximately 230 ns), the memory latency in CPU clock periods is reduced with a lower CPU clock frequency.
- No conflicts of any type are included in the above calculations.
- The numbers indicated are approximate and are the current best estimate.

The CRAY SV1 instruction issue is similar to that of the CRAY J90 series systems. Vector instructions are dispatched from the CIP issue register to the vector issue control. This control can hold five vector instructions. This vector control issues instructions from the vector issue register (VIR) when all vector conflicts have been resolved. The CIP to VIR time is 2 cycles. Vector instruction execution may be 2 cycles faster for instructions that are already waiting in the VIR as compared to the CIP (with no other conflicts or data delays).

Reader Comment Form

Title: CRAY SV1™ Series and CRAY J90se™ Series Programming Differences **Number: 108-0226-001**

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this document?

_____Troubleshooting
_____Reference information
_____Other - please explain
_____Tutorial or introduction
_____Classroom use

_____

Using a scale from 1 (poor) to 10 (excellent), please rate this document on the following criteria and explain your ratings:

_____Accuracy _____
_____Organization _____
_____Readability _____
_____Physical qualities (binding, printing, page layout) _____
_____Amount of diagrams and photos _____
_____Quality of diagrams and photos _____

Completeness (Check one and explain your answer)

_____Too much information _____ Too little information _____Correct amount

_____

_____

_____

You may write additional comments in the space below. Mail your comments to the address below, fax them to us at +1 715 726 4353, or e-mail them to us at *spt@sgi.com.* When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

_____

_____

_____

NAME _____
JOB TITLE _____
E-MAIL ADDRESS _____
SITE/LOCATION _____
TELEPHONE _____
DATE _____

[or attach your business card]

SiliconGraphics
Computer Systems

Attn: Service Publications and Training
890 Industrial Boulevard
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA