

CHAPTER 1 GENERAL

1.1 Introduction

1

The μ PD30412, 30412L (Vr4400MC) processor supports interfaces to secondary cache, system interface, and boot time mode control. This document describes the connection and operation of each of these interfaces.

Remark Please refer also to the following documents when you use this manual.

Vr4000™, Vr4400™ USER'S MANUAL ARCHITECTURE (Document number IEU-1344)

Vr4000PC™, Vr4400PC™ USER'S MANUAL HARDWARE (Document number IEU-1329)

Vr4000SC™, Vr4400SC™ USER'S MANUAL HARDWARE (Document number IEU-1331)

1.2 Operation Fundamentals

A **word** is the basic data element of the Vr4400MC processor. A word is a thirty-two bit data element. A sixty-four bit data element is referred to as a **double word**, a sixteen bit data element is referred to as a **half word** and an eight bit data element is referred to as a **byte**.

1.3 Clocking Fundamentals

The Vr4400MC processor bases all clocking methodology on the single clock input **MasterClock** at the desired operational frequency for the processor. **MasterClock** is multiplied by two internally, using phase locked loop techniques, to generate the processor internal clock, **PClock**. **PClock** is used by the processor's execution units, and to sequence the secondary cache interface. All secondary cache interface transaction protocol and parameters are specified in terms of **PCycles**, where a **PCycle** is the period of **PClock** or half the period of **MasterClock**.

PClock is divided by a programmable divisor to generate the processor internal clock, **SClock**, and the system interface clocks, **TClock** and **RClock**. **SClock** is used by the processor to clock all internal registers that sample system interface inputs and drive system interface outputs. **TClock** and **RClock** are driven off the processor for use by an external agent. The **PClock** to **SClock** divisor is programmed via the boot time mode control interface as described in the Boot Time Mode Control Interface section. All system interface transaction protocol and parameters are specified in terms of **SCycles**, where a **SCycle** is the period of **SClock**, unless otherwise specified.

See CHAPTER 11 CLOCKING for further details on the clocking behavior of the Vr4400MC processor.

★ 1.4 Ordering Information

Part Number	Package	Max. Operating Frequency (MHz) (Internal/External)	Supply Voltage (V)
μ PD30412RJ-50	447-pin ceramic PGA (metal sealed)	100/50	5
μ PD30412RJ-67	447-pin ceramic PGA (metal sealed)	133/67	5
μ PD30412RJ-75	447-pin ceramic PGA (metal sealed)	150/75	5
μ PD30412LRJ-75	447-pin ceramic PGA (seam welded)	150/75	3.3
μ PD30412LRJ-200	447-pin ceramic PGA (seam welded)	200/100	3.45
μ PD30412LRJ-250	447-pin ceramic PGA (seam welded)	250/125	3.45
μ PD30412RP-50	447-pin ceramic PGA (with metal slug)	100/50	5
μ PD30412RP-67	447-pin ceramic PGA (with metal slug)	133/67	5
μ PD30412RP-75	447-pin ceramic PGA (with metal slug)	150/75	5
μ PD30412LRP-75	447-pin ceramic PGA (with metal slug)	150/75	3.3
μ PD30412LRP-200	447-pin ceramic PGA (with metal slug)	200/100	3.45
μ PD30412LRP-250	447-pin ceramic PGA (with metal slug)	250/125	3.45

1.5 Pin Configuration

Fig. 1-1 Pin Configuration (Bottom View)




Table 1-1 Pin Configuration (1/3)

No.	Name	No.	Name	No.	Name	No.	Name
AW37	ColdReset	AA3	SCAddr11	G25	SCData10	AF4	SCData48
AV2	ExtRqst	W3	SCAddr12	E29	SCData11	AJ3	SCData49
C39	Fault	Y6	SCAddr13	G31	SCData12	AJ7	SCData50
AV24	Open <i>Note</i>	W5	SCAddr14	C35	SCData13	AP8	SCData51
AV20	V _{DD}	W7	SCAddr15	K36	SCData14	AT10	SCData52
AV32	IOIn	W1	SCAddr16	N35	SCData15	AR13	SCData53
AV28	IOOut	U3	SCAddr17	AE3	SCData16	AR15	SCData54
AL1	Int0	AN7	SCAddr0W	AG5	SCData17	AT18	SCData55
AA35	IvdAck	AN5	SCAddr0X	AK4	SCData18	AU23	SCData56
AA39	IvdErr	AM6	SCAddr0Y	AN9	SCData19	AT26	SCData57
U39	JTCK	AL7	SCAddr0Z	AU9	SCData20	AR27	SCData58
N39	JTDI	M6	SCDCS	AN13	SCData21	AN29	SCData59
J39	JTDO	G19	SCDChk0	AT14	SCData22	AP32	SCData60
G37	JTMS	T34	SCDChk1	AR17	SCData23	AN35	SCData61
AA37	MasterClock	AP20	SCDChk2	AT22	SCData24	AJ35	SCData62
AJ39	MasterOut	AD34	SCDChk3	AU25	SCData25	AE33	SCData63
B8	ModeClock	C19	SCDChk4	AN27	SCData26	V4	SCData64
AV8	ModeIn	R37	SCDChk5	AR29	SCData27	R5	SCData65
AV16	NMI	AU19	SCDChk6	AN31	SCData28	N5	SCData66
AM34	RClock0	AE37	SCDChk7	AR35	SCData29	E5	SCData67
AL33	RClock1	C17	SCDChk8	AK36	SCData30	G9	SCData68
AW7	RdRdy	N37	SCDChk9	AG35	SCData31	E11	SCData69
AV12	Release	AU17	SCDChk10	T6	SCData32	G13	SCData70
AU39	Reset	AG37	SCDChk11	L3	SCData33	D14	SCData71
Y2	Open <i>Note</i>	E19	SCDChk12	L7	SCData34	C21	SCData72
U5	SCAPar0	R35	SCDChk13	E7	SCData35	D22	SCData73
U1	SCAPar1	AR19	SCDChk14	G11	SCData36	E25	SCData74
P4	SCAPar2	AE35	SCDChk15	E13	SCData37	G27	SCData75
AL5	SCAddr1	R3	SCData0	E15	SCData38	C31	SCData76
AG1	SCAddr2	R7	SCData1	G17	SCData39	F32	SCData77
AE7	SCAddr3	L5	SCData2	C23	SCData40	J35	SCData78
AC1	SCAddr4	F8	SCData3	F24	SCData41	M34	SCData79
AC5	SCAddr5	C9	SCData4	E27	SCData42	AC7	SCData80
AC3	SCAddr6	F12	SCData5	D30	SCData43	AE5	SCData81
AA1	SCAddr7	G15	SCData6	C33	SCData44	AG7	SCData82
AB4	SCAddr8	E17	SCData7	E35	SCData45	AR5	SCData83
AA5	SCAddr9	G21	SCData8	L35	SCData46	AR9	SCData84
AA7	SCAddr10	C25	SCData9	R33	SCData47	AR11	SCData85

Note Leave unconnected.

Remark See CHAPTER 15 PIN SUMMARY for pin functions.

Table 1-1 Pin Configuration (2/3)

No.	Name	No.	Name	No.	Name	No.	Name
AN15	SCData86	AL37	SCData126	U33	Status0	AM38	SysAD30
AP16	SCData87	AG33	SCData127	U35	Status1	AH38	SysAD31
AU21	SCData88	N1	SCOE	V36	Status2	R1	SysAD32
AN23	SCData89	J1	SCTCS	W35	Status3	L1	SysAD33
AR25	SCData90	AN21	SCTChk0	W37	Status4	H2	SysAD34
AP28	SCData91	AN19	SCTChk1	AC37	Status5	E1	SysAD35
AU31	SCData92	AU15	SCTChk2	AC35	Status6	C3	SysAD36
AR33	SCData93	AP12	SCTChk3	AC33	Status7	A5	SysAD37
AL35	SCData94	AU7	SCTChk4	W39	SyncIn	A11	SysAD38
AH34	SCData95	AR7	SCTChk5	AN39	SyncOut	A15	SysAD39
U7	SCData96	AH6	SCTChk6	T2	SysAD0	A23	SysAD40
N3	SCData97	K4	SCTag0	M2	SysAD1	A27	SysAD41
N7	SCData98	G7	SCTag1	J3	SysAD2	A31	SysAD42
C5	SCData99	C7	SCTag2	G3	SysAD3	A35	SysAD43
E9	SCData100	D10	SCTag3	C1	SysAD4	C37	SysAD44
C11	SCData101	C15	SCTag4	A3	SysAD5	E39	SysAD45
C13	SCData102	D18	SCTag5	A9	SysAD6	H38	SysAD46
F16	SCData103	F20	SCTag6	A13	SysAD7	M38	SysAD47
E21	SCData104	E23	SCTag7	A21	SysAD8	AE1	SysAD48
G23	SCData105	D26	SCTag8	A25	SysAD9	AJ1	SysAD49
C27	SCData106	C29	SCTag9	A29	SysAD10	AM2	SysAD50
F28	SCData107	G29	SCTag10	A33	SysAD11	AR1	SysAD51
E31	SCData108	E33	SCTag11	B38	SysAD12	AU3	SysAD52
G33	SCData109	G35	SCTag12	E37	SysAD13	AW5	SysAD53
J37	SCData110	L33	SCTag13	G39	SysAD14	AW11	SysAD54
N33	SCData111	L37	SCTag14	L39	SysAD15	AW15	SysAD55
AD6	SCData112	P36	SCTag15	AD2	SysAD16	AW23	SysAD56
AG3	SCData113	AF36	SCTag16	AH2	SysAD17	AW27	SysAD57
AJ5	SCData114	AJ37	SCTag17	AL3	SysAD18	AW31	SysAD58
AU5	SCData115	AJ33	SCTag18	AN3	SysAD19	AW35	SysAD59
AN11	SCData116	AN37	SCTag19	AU1	SysAD20	AU37	SysAD60
AU11	SCData117	AU35	SCTag20	AW3	SysAD21	AR39	SysAD61
AU13	SCData118	AR31	SCTag21	AW9	SysAD22	AL39	SysAD62
AN17	SCData119	AU29	SCTag22	AW13	SysAD23	AG39	SysAD63
AR21	SCData120	AN25	SCTag23	AW21	SysAD24	A17	SysADC0
AP24	SCData121	AR23	SCTag24	AW25	SysAD25	R39	SysADC1
AU27	SCData122	J5	SCWrW	AW29	SysAD26	AW17	SysADC2
AT30	SCData123	J7	SCWrX	AW33	SysAD27	AD38	SysADC3
AU33	SCData124	H6	SCWrY	AV38	SysAD28	A19	SysADC4
AN33	SCData125	G5	SCWrZ	AR37	SysAD29	T38	SysADC5


Remark See CHAPTER 15 PIN SUMMARY for pin functions.

Table 1-1 Pin Configuration (3/3)

No.	Name	No.	Name	No.	Name	No.	Name
AW19	SysADC6	H36	V _{DD}	B22	Gnd	AP30	Gnd
AC39	SysADC7	K6	V _{DD}	B30	Gnd	AP34	Gnd
G1	SysCmd0	K38	V _{DD}	B36	Gnd	AP36	Gnd
E3	SysCmd1	P2	V _{DD}	D2	Gnd	AT2	Gnd
B2	SysCmd2	P34	V _{DD}	D6	Gnd	AT6	Gnd
B12	SysCmd3	T4	V _{DD}	D12	Gnd	AT12	Gnd
B16	SysCmd4	T36	V _{DD}	D20	Gnd	AT20	Gnd
B20	SysCmd5	V6	V _{DD}	D28	Gnd	AT28	Gnd
B24	SysCmd6	V38	V _{DD}	D34	Gnd	AT34	Gnd
B28	SysCmd7	Y38	V _{DD}	D38	Gnd	AT38	Gnd
B32	SysCmd8	AB2	V _{DD}	F4	Gnd	AV4	Gnd
A37	SysCmdP	AB34	V _{DD}	F6	Gnd	AV10	Gnd
H34	TClock0	AD4	V _{DD}	F10	Gnd	AV18	Gnd
J33	TClock1	AD36	V _{DD}	F18	Gnd	AV26	Gnd
AE39	V _{DD} Ok	AF6	V _{DD}	F26	Gnd	AV36	Gnd
AN1	ValidIn	AF38	V _{DD}	F34	Gnd		
AR3	ValidOut	AK2	V _{DD}	F36	Gnd		
A7	WrRdy	AK34	V _{DD}	K2	Gnd		
W33	V _{DD} Sense	AM4	V _{DD}	K34	Gnd		
U37	GndSense	AM36	V _{DD}	M4	Gnd		
AA33	V _{DD} P	AP2	V _{DD}	M36	Gnd		
Y34	GndP	AP10	V _{DD}	P6	Gnd		
A39	V _{DD}	AP18	V _{DD}	P38	Gnd		
B6	V _{DD}	AP26	V _{DD}	V2	Gnd		
B10	V _{DD}	AP38	V _{DD}	V34	Gnd		
B18	V _{DD}	AT4	V _{DD}	Y4	Gnd		
B26	V _{DD}	AT8	V _{DD}	Y36	Gnd		
B34	V _{DD}	AT16	V _{DD}	AB6	Gnd		
D4	V _{DD}	AT24	V _{DD}	AB36	Gnd		
D8	V _{DD}	AT32	V _{DD}	AB38	Gnd		
D16	V _{DD}	AT36	V _{DD}	AF2	Gnd		
D24	V _{DD}	AV6	V _{DD}	AF34	Gnd		
D32	V _{DD}	AV14	V _{DD}	AH4	Gnd		
D36	V _{DD}	AV22	V _{DD}	AH36	Gnd		
F2	V _{DD}	AV30	V _{DD}	AK6	Gnd		
F14	V _{DD}	AV34	V _{DD}	AK38	Gnd		
F22	V _{DD}	AW1	V _{DD}	AP4	Gnd		
F30	V _{DD}	AW39	V _{DD}	AP6	Gnd		
F38	V _{DD}	B4	Gnd	AP14	Gnd		
H4	V _{DD}	B14	Gnd	AP22	Gnd		

Remark See CHAPTER 15 PIN SUMMARY for pin functions.

Fig. 1-2 Pad Configuration (Top View)



Caution The voltage of the metal seal on the top of the package is equal to V_{DD} .

Table 1-2 Pad Configuration


No.	Name	No.	Name	No.	Name	No.	Name
1	V_{DD}	7	V_{DDP}	13	V_{DD}	19	V_{DD}
2	Gnd	8	PLLCap1	14	Gnd	20	Gnd
3	PLLCap0	9	V_{DD}	15	V_{DD}	21	V_{DD}
4	GndP	10	Gnd	16	Gnd	22	Gnd
5	GndP	11	V_{DD}	17	V_{DD}		
6	V_{DDP}	12	Gnd	18	Gnd		

Table 1-3 Capacitance of Chip Capacitors

Pads	Capacitance [μF]	Pads	Capacitance [μF]
1 – 2	0.1	13 – 14	0.1
3 – 4	0.001	15 – 16	0.1
5 – 6	0.1	17 – 18	0.1
7 – 8	0.001	19 – 20	0.1
9 – 10	0.1	21 – 22	0.1
11 – 12	0.1		

Remark Chip capacitors are incorporated at delivery.

Fig. 1-3 Package Side View



CHAPTER 2 CACHE COHERENCY

The VR4400MC processor manages its primary and secondary caches using a write back methodology, that is, stores write data into the caches, but a modified cache line is not written back to memory until the cache line is replaced, or until the cache line is exported or flushed from the secondary cache. When the contents of a cache line is not consistent with memory, it is said to be dirty. Many systems, in particular multiprocessor systems, or systems that employ input/output (IO) devices that are capable of direct memory access (DMA), may require the system to behave as if the caches are always consistent with memory and each other. Schemes for maintaining consistency between multiple write back caches or between write back caches and memory are referred to as **cache coherency protocols**.

The processor, in its secondary cache mode, provides a set of cache states and mechanisms for manipulating the contents and state of the cache that are sufficient to implement a variety of cache coherency protocols, both snoopy and directory based. In particular, the processor supports both the write invalidate and write update protocols simultaneously.

The coherency protocol for lines in the cache is controllable via bits in the translation look-aside buffer (TLB) on a per TLB page basis. Specifically, the TLB contains three bits per entry that control the coherency attributes of a page. The three bits are encoded to provide five possible coherency attributes per page, uncached, cacheable noncoherent, cacheable coherent exclusive (exclusive), cacheable coherent exclusive on write (sharable), and cacheable coherent update on write (update). A processor in the no-secondary cache mode supports only the uncached and noncoherent coherency attributes. The supported page attributes depend on the configuration of the processor as illustrated below.

Table 2-1 Allowed Page Attributes

<u>Configuration</u>	<u>Uncached</u>	<u>Cacheable Noncoherent</u>	<u>Exclusive</u>	<u>Sharable</u>	<u>Update</u>
VR4000PC, VR4400PC	Supported	Supported	NA	NA	NA
VR4000SC, VR4400SC	Supported	Supported	Supported	NA	NA
VR4400MC	Supported	Supported	Supported	Supported	Supported
Remark NA: Not Available					

If a page has the coherency attribute uncached, the processor will issue a word or partial word read or write directly to main memory for any load or store to a location within that page. Lines within an uncached page are assumed never to be cache resident.

If the coherency attribute is sharable, the processor will issue a coherent block read for a load miss to a location within the page, and a coherent block read that requests exclusivity for a store miss to a location within the page. In most systems, coherent reads require snoops or directory checks to occur while noncoherent reads do not. A coherent read that requests exclusivity implies that the processor will function most efficiently if the requested cache line is returned to it in an exclusive state, but the processor will still perform correctly if the cache line is returned in a shared state. Cache lines within the page will be managed with a write invalidate protocol, that is the processor will issue an invalidate on a store hit to a shared cache line.

If the coherency attribute is update, the processor will issue a coherent block read for a load or store miss to a location within the page. Cache lines within the page will be managed with a write update protocol, that is the processor will issue an update on a store hit to a shared cache line.

If the coherency attribute is exclusive, the processor will issue a coherent block read that requests with a write invalidate protocol. Load linked store conditional instruction sequences must insure that the link location is not in a page managed with the exclusive coherency attribute.

If the coherency attribute is noncoherent, the processor will issue a noncoherent block read for a load or store miss to a location within the page.

The encoding of the coherency attributes in the TLB is specified in V_R4000, V_R4400 User's Manual Architecture. The behavior of the processor on load misses, store misses, and store hits to shared cache lines for each of the coherency attributes is summarized in Table 2-2 Coherency Attributes and Processor Behavior.

Table 2-2 Coherency Attributes and Processor Behavior

<u>Attribute</u>	<u>Load Miss</u>	<u>Store Miss</u>	<u>Store Hit Shared</u>
uncached	Main memory read	Main memory write	NA
noncoherent	Noncoherent read	Noncoherent write	Invalidate ^{Note}
exclusive	Coherent read exclusive	Coherent read exclusive	Invalidate ^{Note}
sharable	Coherent read	Coherent read exclusive	Invalidate
update	Coherent read	Coherent read	Update
Note This should not occur under normal circumstances.			
Remark NA: Not Available			

The following sections describe the primary and secondary cache states provided by the processor, the cache state transitions performed by the processor during execution, and the mechanisms provided for an external agent to manipulate the state and contents of the primary and secondary cache.

2.1 Cache States

The V_R4400MC maintains four primary cache states and five secondary cache states. The five secondary cache states are:

- Invalid
- Shared
- Dirty shared
- Clean Exclusive
- Dirty Exclusive

The four primary cache states are:

- Invalid
- Shared
- Clean Exclusive
- Dirty Exclusive

The primary cache state shared corresponds to the secondary cache states shared and dirty shared.

The cache states and line attributes are illustrated below.

Table 2-3 Primary and Secondary Cache States and Line Attributes
(a) Primary Cache

<u>Cache State</u>	<u>Line Attributes</u>
Invalid	The cache line does not contain valid information.
Valid	The cache line contains valid information (primary instruction cache only).
Shared	The cache line contains valid information and may be present in another processor's cache. The cache line may or may not be consistent with memory, and may or may not be owned.
Clean Exclusive (CE)	The cache line contains valid information and is not present in any other processor's cache. The cache line is consistent with memory and owned.
Dirty Exclusive (DE)	The cache line contains valid information and is not present in any other processor's cache. The cache line is inconsistent with memory and owned.

(b) Secondary Cache

<u>Cache State</u>	<u>Line Attributes</u>
Invalid	The cache line does not contain valid information.
Shared	The cache line contains valid information and may be present in another processor's cache. The cache line may or may not be consistent with memory, and is not owned.
Dirty Shared	The cache line contains valid information and may be present in another processor's cache. The cache line is inconsistent with memory and owned.
Clean Exclusive (CE)	The cache line contains valid information and is not present in any other processor's cache. The cache line is consistent with memory and owned.
Dirty Exclusive (DE)	The cache line contains valid information and is not present in any other processor's cache. The cache line is inconsistent with memory and owned.

The cache state of a line in the processor's primary or secondary cache indicates the validity, shared, dirty and ownership attributes of the cache line. A cache line that does not contain valid information must be marked **invalid**; a cache line in any state other than invalid contains valid information. A cache line that is present in more than one cache in the system is said to be **shared** and must be in one of the shared states. A cache line that is present in exactly one cache in the system is said to be **exclusive** and may be in one of the exclusive states. A cache line that contains data that is consistent with memory is said to be **clean** and may be in one of the clean states. A cache line that contains data that is not consistent with memory is said to be **dirty** and must be in one of the dirty states, or in the shared state. The processor has a concept of ownership for cache lines. When the processor is the owner of a particular cache line it is responsible for writing the cache line back to memory when it is replaced in the course of satisfying a cache miss, or during the execution of a cache instruction. A cache line is owned by the processor if its secondary cache state is dirty exclusive or dirty shared. Note that the cache states have no distinction between clean and dirty in the primary instruction cache, since the processor does not use a write back methodology in managing the primary instruction cache.

The primary and secondary cache states have been chosen to maintain all of the state information that the processor may need during execution in the primary cache, while maintaining all of the state information that an external agent may need to manage a cache coherency protocol in the secondary cache.

The allowed cache states for the primary cache are illustrated below.

Table 2-4 Allowed Primary and Secondary Cache States

	<u>Vr4000PC, Vr4400PC</u>		<u>Vr4000SC, Vr4400SC</u>		<u>Vr4400MC</u>	
	<u>Instruction Cache</u>	<u>Data Cache</u>	<u>Instruction Cache</u>	<u>Data Cache</u>	<u>Instruction Cache</u>	<u>Data Cache</u>
Primary Cache	Invalid Valid	Invalid DE Note 1	Invalid Valid Note 2	Invalid CE DE	Invalid Valid Note 2	Invalid Shared CE DE
Secondary Cache		NA		Invalid CE DE		Invalid Shared Dirty Shared CE DE
Notes 1. Used with the W bit of the primary cache tag to indicate whether the current line is modified and must be written back to memory or secondary cache (W = 0 clean; W = 1 modified). 2. The Valid state of the Primary Instruction Cache is mapped to any valid Secondary Cache states (Clean Exclusive, Dirty Exclusive, Shared, and Dirty Shared). Remark DE: Dirty Exclusive, CE: Clean Exclusive, NA: Not Available						

2.2 Cache State Changes During Processor Execution

The initial state of a cache line is specified by an external agent when it supplies the cache line. During the course of processor execution, the processor may change the state of a cache line. The following events will cause changes to the state of the cache:

A store to the primary cache will write the data in the primary cache and change the W bit to one. The processor will also change the secondary cache state from a clean state to a dirty state without changing the data. The processor will not access the secondary cache on any subsequent stores to the same cache line in the primary cache.

A store to a shared cache line, that is a line marked shared in the primary cache and either shared or dirty shared in the secondary cache, will cause the processor to issue either an invalidate request or an update request depending on the coherency attribute in the TLB entry for the page that contains the cache line. Upon successful completion of an invalidate, the processor will complete the store and change the state of the cache line to dirty exclusive in both the primary and secondary caches. Upon successful completion of an update, the processor will complete the store and change the state of the cache line to shared in the primary cache and dirty shared in the secondary cache if dirty shared mode (ModeBit [3]) is enabled. Dirty shared mode is programmable via the boot time mode control interface. If dirty shared mode is not enabled, the state of the primary and secondary caches will be left unchanged after successful completion of an update.

2.3 Cache Line Write Back

The processor will write a cache line back to memory when it is replaced, or written back to memory as the result of executing a cache instruction, if the cache line is in the state dirty exclusive or dirty shared in the secondary cache. When the processor writes a cache line back to memory, it does not ordinarily retain a copy of the cache line, and the state of the cache line is changed to invalid. However, if a cache line is written back to memory using the hit writeback cache instruction, the processor will retain a copy of the cache line. If the cache line is retained, the processor will change its state to clean exclusive if the secondary cache state was dirty exclusive before the write or shared if the secondary cache state was dirty shared before the write.

Whether or not the processor is retaining the line is signaled by the processor during a write.

2.4 Manipulation of the Caches by an External Agent

The Vx4400MC provides mechanisms for an external agent to examine and manipulate the state and contents of the primary and secondary caches:

An external agent must specify the state in which data, supplied in response to a processor read request, is to be loaded into the processor's caches. Data may be loaded in any of the four valid secondary cache states. Data returned by the external agent must not be marked as invalid. The secondary cache state will be mapped to a primary cache state as described previously.

An external agent may issue a snoop request to the processor which will cause the processor to return the secondary cache state of the specified cache line. At the same time it will change the state of the specified cache line in both the primary and secondary caches, according to a state change function specified by the external agent, atomically with respect to the response to the snoop request.

An external agent may issue an invalidate request or an update request to the processor. An invalidate request will cause the processor to change the state of the specified cache line to invalid in both the primary and secondary caches. An update request will cause the processor to write the specified data element into the specified cache line, and either change the state of the cache line to shared in both the primary and secondary caches, or leave the state of the cache line unchanged, depending on the nature of the update request. An external agent may issue updates, without changing the state of the cache line, to cache lines that are either in exclusive or shared states. If an update request is issued to the primary instruction cache, the secondary cache line and primary instruction cache line are updated.

An external agent may issue an intervention request which will cause the processor to return the secondary cache state of the specified cache line, and the contents of the specified secondary cache line under certain conditions related to the state of the cache line and the nature of the intervention request. At the same time the processor will change the state of the specified cache line in both the primary and secondary caches, according to a state change function specified by the external agent, atomically with respect to the response to the intervention request.

2.5 Cache Line Ownership

The V_r4400MC has a concept of ownership for cache lines. The ownership of a cache line is maintained as follows:

A processor assumes ownership of a cache line when the state of the cache line transitions to dirty shared or dirty exclusive. For responses to processor coherent read requests in which the data is returned with an indication that it must be loaded in the dirty shared or dirty exclusive state, the cache state is set at the completion of the read response when the last word of read response data is returned. Therefore, the processor will assume ownership of the cache line when the last word of read response data is returned.

The processor gives up ownership of a cache line when the state of the cache line transitions to invalid, shared, or clean exclusive. For processor write requests the state of the cache line will transition to invalid if the cache line is replaced, or clean exclusive or shared if the cache line is retained. In either case, the cache state transition will occur at the completion of the write request when the last word of write data is transmitted to the external agent. Therefore, the processor will give up ownership of the cache line when the last word of write data is transmitted to the external agent.

For external requests, other than read responses, any cache state change associated with the external request will occur at the completion of the external request and therefore any change of ownership resulting from the cache state change will occur at the completion of the external request.

2.6 Ordering Considerations

Many cache coherent multiprocessor systems must obey ordering constraints on stores to shared data such that they exhibit the same behavior as a uniprocessor system in a multi-programming environment. A multiprocessor system that exhibits such behavior is said to be **strongly ordered**.

An algorithm typically used to test for strong ordering is the following: Processor A does a store to location X at the same time processor B does a store to location Y. Locations X and Y have no particular relationship, i.e. they are not in the same cache line. Next processor A does a load from location Y at the same time that processor B does a load from location X. In order for the system to be strongly ordered either processor A must load the new value of Y, or processor B must load the new value of X, or both processors A and B must load the new values of Y and X respectively under all conditions. If both processors A and B load the old values of Y and X, under any conditions, the system does not meet the requirements for strong ordering. The algorithm to test for strong ordering is summarized below.

<u>Processor A</u>	<u>Processor B</u>
Store to location X	Store to location Y
Load from location Y	Load from location X

In order for the above test algorithm for strong ordering to succeed stores must have a global ordering in time; that is, every processor in the system must agree that either the store to location X preceded the store to location Y, or the store to location Y preceded the store to location X. If this global ordering is enforced the above test algorithm for strong ordering will succeed as described.

In a system employing the V_r4400MC processor the requirements to achieve strong ordering translate to a need to precisely control when the processor restarts after completion of a processor coherence request with respect to cache state changes stemming from external coherence requests. Specifically, a system designer must make sure that any cache state changes, resulting from external coherence requests that occur before a processor coherence request, are completed before the processor is allowed to restart after completion of the processor coherence request.

The Vr4400MC processor obeys the following paradigms for restart after issuing a coherence request:

For coherent read requests, the processor will restart after the requested double word is transmitted to the processor if sub-block ordering is enabled or after the last word in the block is transmitted to the processor if sequential ordering is enabled, unless a processor invalidate or update request is unacknowledged. Any external requests that must be completed before the read is complete must be issued to the processor before the read response is issued.

For write requests, the processor will restart after the write request is complete, that is, after the last double word of data associated with the write request has been transmitted to the external agent unless a processor read request is pending or a processor invalidate or update request is unacknowledged.

For invalidate and update requests, the processor will restart after the assertion of $\overline{\text{IvdAck}}$ or $\overline{\text{IvdErr}}$ unless a processor read request is pending or unless it is processing an external request when $\overline{\text{IvdAck}}$ or $\overline{\text{IvdErr}}$ is asserted. If $\overline{\text{IvdAck}}$ or $\overline{\text{IvdErr}}$ is asserted during or after the first cycle that the external agent asserts $\overline{\text{ExtRqst}}$, the processor will accept the external request and complete any cache state changes associated with it before the processor restarts, otherwise, the processor will restart before beginning the external request. For any external requests that must be completed before a processor invalidate or update completes, the external agent must assert $\overline{\text{ExtRqst}}$ before or during the cycle that signal $\overline{\text{IvdAck}}$ or $\overline{\text{IvdErr}}$ is asserted.

2

CHAPTER 3 SECONDARY CACHE INTERFACE

The V_a4400MC is designed to operate with an external secondary cache. The secondary cache is accessible by the processor and to the system interface. The cache contains data, cache tags and cache line state bits.

3.1 Secondary Cache Overview

The V_a4400MC secondary cache is assumed to consist of one bank of industry standard static RAMs with output enables. The V_a4400MC secondary cache consists of quad-word (128 bit) wide data array and a 25-bit wide tag array. Check fields are added to both the data and tag arrays to improve data integrity. The secondary cache may be configured as joint or split instruction/data. The maximum secondary cache size is 4 Mbytes and the minimum secondary cache size is 128 Kbytes for joint and 256 Kbytes for split instruction/data. The secondary cache is direct-mapped, and is addressed with the lower part of the physical address.

3

3.2 Secondary Cache Interface Signal Description

The signals that connect the V_a4400MC processor to its secondary cache are described in this section.

3.2.1 Secondary Cache Interface Signal Summary

SCData(127:0):	(i/o) A 128-bit bus used to read or write cache data from/to the secondary cache.									
SCDChk(15:0):	(i/o) A 16-bit bus which conveys two ECC fields that cover the upper or lower 64 bits of the SCData from/to the secondary cache.									
SCTag(24:0):	(i/o) A 25-bit bus used to read or write cache tags from/to the secondary cache.									
SCTChk(6:0):	(i/o) A 7-bit bus which conveys an ECC field that covers the SCTag from/to the secondary cache.									
SCAddr(17:1)	(o) A 17-bit bus which addresses the secondary cache.									
SCAddr0Z:	(o) Bit 0 of the secondary cache address.									
SCAddr0Y:	(o) Bit 0 of the secondary cache address.									
SCAddr0X:	(o) Bit 0 of the secondary cache address.									
SCAddr0W:	(o) Bit 0 of the secondary cache address.									
SCAPar(2:0):	(o) The secondary cache address even parity bus cover the following bits: <table><tr><td>SCAPar(2)</td><td>7 bits:</td><td><u>SCWr</u>, SCAddr(17:12)</td></tr><tr><td>SCAPar(1)</td><td>7 bits:</td><td><u>SCDCS</u>, SCAddr(11:6)</td></tr><tr><td>SCAPar(0)</td><td>7 bits:</td><td><u>SCTCS</u>, SCAddr(5:0)</td></tr></table>	SCAPar(2)	7 bits:	<u>SCWr</u> , SCAddr(17:12)	SCAPar(1)	7 bits:	<u>SCDCS</u> , SCAddr(11:6)	SCAPar(0)	7 bits:	<u>SCTCS</u> , SCAddr(5:0)
SCAPar(2)	7 bits:	<u>SCWr</u> , SCAddr(17:12)								
SCAPar(1)	7 bits:	<u>SCDCS</u> , SCAddr(11:6)								
SCAPar(0)	7 bits:	<u>SCTCS</u> , SCAddr(5:0)								
<u>SCOE</u>:	(o) A signal which enables the outputs of the secondary cache RAMs.									
<u>SCWrZ</u>:	(o) Secondary cache write enable.									
<u>SCWrY</u>:	(o) Secondary cache write enable.									
<u>SCWrX</u>:	(o) Secondary cache write enable.									
<u>SCWrW</u>:	(o) Secondary cache write enable.									
<u>SCDCS</u>:	(o) A signal which enables the chip select pins of the secondary cache RAMs associated with SCData and SCDChk.									
<u>SCTCS</u>:	(o) A signal which enables the chip select pins of the secondary cache RAMs associated with SCTag and SCTChk.									

3.2.2 Details of Secondary Cache Interface Signals

The interface to the V_R4400MC secondary cache is designed to maximize the efficiency of servicing primary cache misses. The width of the data portion of secondary cache interface is chosen to be 128 bits to support a data rate into the primary cache that is near the processor to primary cache bandwidth during normal operation. To assure that this bandwidth is maintained, each data, tag and check pin must be connected to only one static RAM device. The **SCAddr** bus, the **SCOE** signal, the **SCDCS** signal and the **SCTCS** drive a large number of static RAM devices, so one level of external buffering between the V_R4400MC and the cache array is necessary.

The speed of the secondary cache interface is limited by buffered control signals. Critical control signals are duplicated to minimize this effect. The **SCWr** signal and **SCAddr(0)** are duplicated four times so that external buffering will not be required. When an 8-word (256-bit) primary cache line is used, these signals can be controlled significantly faster to reduce the time of the two back-to-back transfers. These duplicated control signals are specified to drive 11 parts each, so that a total of 44 RAM packages can be used in the cache array. This permits a cache design using 64 Kbyte by 4 bit or 256 Kbyte by 4 bit standard static RAMs. Other cache designs are also acceptable, for example a smaller cache design using 228 Kbyte by 8 bit static RAMs as it would present less load on the address pins and control signals. Note that duplicated signals like **SCWrW**, **SCWrX**, **SCWrY** and **SCWrZ** will be described in this document as though they were a single signal, which in this case is called **SCWr**.

The benefit of duplicating **SCAddr(0)** will be greater if fast sequential static cache RAMs become available. If **SCAddr(0)** is attached to a static RAM address bit that effects column decode only, the read cycle time with respect to that pin should approximate the output enable time of the RAM and for fast static RAMs should be half that of the nominal read cycle time.


When the split instruction/data cache mode is enabled, assertion of the top **SCAddr** bit, **SCAddr(17)** will enable the instruction half of the cache instead of the data half.

It is possible to design a cache that supports both joint and split instruction/data configurations with less than the maximum cache size. **SCAddr(12:0)** must be used to address the cache in all configurations. **SCAddr(17)** must be used to support the split instruction/data configuration. Any of **SCAddr(16:13)** may be omitted because of the fixed width of the physical tag array.

The **SCDChk** bus is divided into two fields to cover the upper and lower 64 bits of **SCData**. This form is required to keep the width of internal data paths to 64 bits.

The **SCTag** bus is divided into three fields, as shown in Fig. 3-1 **SCTag Fields**. The lower 19 bits consist of the upper physical address bits. The upper three bits consist of cache state, which can be one of Invalid (I), Clean Exclusive (CE), Dirty Exclusive (DE). Bits 19 – 21 are used to maintain information about the virtual address used for caching parts of a secondary cache line in the primary cache. This field holds bits 14 through 12 of the virtual address.

Fig. 3-1 SCTag Fields



Bits 19 – 21 of **SCTag** are needed to locate entries in the primary caches. The V_R4400MC has two primary caches, one for instruction and one for data, which are direct-mapped and are indexed using a subset of the lower 15 bits of the virtual address (implementation dependent, based on primary cache size). If a cache coherency request is processed that requires a cache state change or invalidation, the middle three bits of the **SCTag** portion of the secondary cache allow primary cache lines affected by that cache coherency request to be found. The three bits of information stored are bits 14 through 12 of the virtual address. This information is loaded during secondary cache misses. On each secondary cache access the virtual address bits are compared with the values found in the secondary

cache tag. If a mismatch occurs, a trap is taken and the trap handler can modify the bits in the secondary cache tag to hold the new values, and the old values are used by the trap handler to purge primary cache locations, so that all primary cache lines holding valid data have indexes known to the secondary cache. This mechanism also helps preserve the integrity of cached accesses to a physical address using differing virtual addresses known as virtual synonyms.

The $\overline{\text{SCDCS}}$ and $\overline{\text{SCTCS}}$ are needed to disable reads or writes of the data array or tag array when the other array is being accessed. These signals are useful for saving power on snoop and invalidate requests, as accesses to the data array are not necessary. These signals are also useful for writing data from the data primary cache to the secondary cache, as the secondary cache state cannot always be determined from the primary cache state.

3

3.3 Control of Secondary Cache Interface

The control of the secondary cache is configurable for various clock rates and static RAM speeds. All configurable parameters are specified in multiples of PClock, which runs at twice the frequency of the external system clock, MasterClock. Boot time mode control registers will hold the various configuration parameters, so that they can be specified when initializing the VR4400MC. The table below shows the number of PClock cycles that those parameters can be specified.

Table 3-1 Secondary Cache Read/Write Cycle Parameters

Parameter	Settable Number of PCycles
tRd1Cyc	4 – 15
tRd2Cyc	3 – 15
tDis	2 – 7
tWr1Dly	1 – 3
tWr2Dly	1 – 3
tWrRC	0 – 1
tWrSUP	3 – 15


3.3.1 Read Cycles

Each secondary cache read sequence begins with the driving of the address pins. The output enable signal $\overline{\text{SCOE}}$ is asserted at the same time.

There are two basic read cycles: a four-word read, and an eight-word read.


For the four-word read, there are two parameters of interest. The first parameter is read sequence cycle time, tRd1Cyc, which specifies the time from the driving of the SCAddr bus to the sampling of the SCData bus. The second parameter is the cache output disable time tDis, which specifies the time from the end of a read cycle to the start of the next write cycle. Fig. 3-2 Four-Word Read Cycle illustrates the four-word read sequence.

Fig. 3-2 Four-Word Read Cycle



For the eight-word read, there is one additional parameter of interest: the time from the first sample point to the second sample point, TRd2Cyc. The lower order address bit, SCAddr(0) is changed at the same time as the first read sample point. Fig. 3-3 Eight-Word Read Cycle illustrates the eight-word read sequence.

Fig. 3-3 Eight-Word Read Cycle



All read cycles can be aborted by changing the address. A new cycle starts beginning with the edge on which the address is changed. Additionally, the period TDis after a read cycle can be interrupted any time by the start of a new read cycle. If a read cycle is aborted by a write cycle, $\overline{SCOE}$ must be deasserted for the TDis period, before the write cycle can commence. Read cycles can also be extended indefinitely. There is no requirement to change the address at the end of a read cycle.

3.3.2 Write Cycles


Like the read sequence, the secondary cache write sequence begins with the driving of the address pins.

There are two basic write cycles: a four-word write, and an eight-word write.

For the four-word write, there are several parameters of interest. The first parameter, T_{Wr1Dly} is the time from driving address to the assertion of $\overline{SCWr}$. The second parameter, T_{WrSUP} is the time from driving the second data double-word to the deassertion of $\overline{SCWr}$. The final parameter, T_{WrRc} , is the time from the deassertion of $\overline{SCWr}$ to the beginning of the next cycle. T_{WrRc} will be zero for most cache designs. Note that the upper data double word and the lower data double word will normally be driven one cycle apart. This reduces the peak current consumption in the output drivers. **Fig. 3-4 Four-Word Write Cycle** illustrates the four-word write sequence. The order of driving the upper versus the lower halves of $SCData$ are not fixed, either the upper or the lower halves might be driven first.


3

Fig. 3-4 Four-Word Write Cycle



The eight-word write has one new parameter. The parameter, T_{Wr2Dly} , is the time from changing the low-order address bit $SCAAddr(0)$ to the assertion of $\overline{SCWr}$ the second time. The lower half of $SCData$ will be driven out on the same edge as the change in $SCAAddr(0)$. **Fig. 3-5 Eight-Word Write Cycle** illustrates the eight-word write sequence.

Fig. 3-5 Eight-Word Write Cycle



Note The secondary cache tag and tag check are identical for both four-word write cycles.

When receiving data from the system interface, it is possible that the first data double word will arrive several cycles before the second. In this case, the cache state machine will simply wait until that data is available before asserting **SCWr** and will extend the **SCWr** until **TWRSup** after the driving of the second data item.

CHAPTER 4 SYSTEM INTERFACE

The system interface allows the processor access to external resources required to satisfy cache misses while also allowing an external agent access to certain processor internal resources. In the V_R4400MC, the system interface also provides the processor mechanisms with which to maintain the cache coherency of shared data, while providing an external agent mechanisms with which to maintain system wide multiprocessor cache coherency.

4.1 System Interface Overview

An event that occurs within the processor that requires access to external system resources will be referred to as a **system event**. System events include, a load that misses in both the primary and secondary caches, a store that misses in both the primary and secondary caches, a store that hits in either the primary or secondary data cache on a shared line, and an uncached load or store. A miss in both caches will require the write back to memory of the cache line that is being replaced if it is in one of the dirty cache states. Cache instructions will also cause system events under certain circumstances. For more details on V_R4400MC cache instructions see V_R4000, V_R4400 USER'S MANUAL—ARCHITECTURE.

When a system event occurs, the processor will issue a request or a series of requests called **processor requests** through the system interface to access some external resource and service the event. The processor's system interface must be connected to some external agent that understands the system interface protocol and can coordinate the access to system resources.

Processor requests include read, write, null write, invalidate, and update requests. Reads are requests for a block, double word, word or partial word of data from main memory or another system resource. Writes provide a block, double word, word or partial word of data to be written to main memory or another system resource. Null writes indicate that an expected write has been obviated as a result of some external request. Invalidates are requests to invalidate the specified cache line in every other cache in the system. Updates are requests to update every other cache in the system with the specified double word, word, or partial word of data.

An external agent may require access to the processor's caches or to some processor internal resource. In this event the external agent will issue a request to the processor through the system interface called an **external request** to provide the access.

External requests include read, write, invalidate, update, snoop, intervention, and null requests. Reads are requests for a word of data from some processor internal resource. Writes provide a word of data to be written to some processor internal resource. Invalidates specify a cache line that must be marked invalid in the processor's primary and secondary caches. Updates provide a double word, word or partial word of data to be written to the processor's primary and secondary caches. Snoop requests are used to interrogate the processor's secondary cache to discover if the processor has a valid copy of a particular cache line and if so what cache state the line is in. Snoop requests require the processor to return an indication of the state of the cache line at the specified physical address in the secondary cache if it is present. Intervention requests require the processor to return an indication of the state of the cache line at the specified physical address in the secondary cache and the contents of the cache line from the primary and secondary caches under certain conditions related to the state of the cache line and the nature of the intervention request. Null requests require no action by the processor, rather they simply provide a mechanism for an external agent to either return control of the secondary cache to the processor, or to return control of the system interface to the processor.

When the processor or an external agent receives a read request, it must access the specified resource and return the requested data. For external read requests, the data will be returned directly in response to the read request. For processor read requests the read request and the return of data by an external agent in response to the read request are disconnected or split. The response data may be returned at any time after the read request, and the system interface is not in use by the read during the time between the read request and the return of response data.

An external agent may initiate an unrelated external request before it returns the response data for a processor read. The return of data in response to a processor read request will be accomplished via a read response. While a read response is technically also an external request, read responses have different characteristics from all other external requests in that arbitration for the system interface must not be performed for read responses. For this reason, read responses will be treated separately from all other external requests, and will be called simply read responses.

Processor read requests that have been issued but for which data has not yet been returned are said to be pending. A read is pending until the read data has been returned. Note that the data identifier associated with the response data may signal that the returned data is erroneous, causing the processor to take a bus error.

External read requests are not split. The system interface is in use between the read request and the return of data by the processor.

A processor read request is complete after the last word of response data has been received from the external agent. A processor write request is complete after the last word of data has been transmitted. A processor invalidate or update request requires a completion acknowledge via the invalidate acknowledge signals $\overline{\text{IvdAck}}$ or $\overline{\text{IvdErr}}$, unless the invalidate or update is canceled by the external agent. Invalidate or update cancellation is signaled to the processor during external invalidate, update, snoop, and intervention requests. External invalidate, update, snoop and intervention requests, as a group are referred to as external coherence requests. $\overline{\text{IvdAck}}$ is used to signal that a processor invalidate or update request has succeeded. $\overline{\text{IvdErr}}$ is used to signal that a processor invalidate or update request has failed. Since the completion acknowledge for processor invalidate and update requests is signaled through the system interface on dedicated pins, the completion acknowledge may occur in parallel with processor and external requests. A processor invalidate or update request that has been issued but for which the processor has not yet received an acknowledge or a cancellation is said to be unacknowledged.

An external read request is complete after the processor returns the requested word of data. An external write request is complete after the word of data has been transmitted. An external invalidate or update request is complete after the request has been transmitted. An external snoop request is complete after the processor returns the state of the specified cache line. An external intervention request is complete after the processor returns the state of the specified cache line, if the processor does not return the contents of the cache line, or after the processor returns the last word of data for the specified cache line.

The processor must manage the flow of processor requests and external requests. The flow of external requests is controlled by the processor via the external request arbitration signals $\overline{\text{ExtRqst}}$, and $\overline{\text{Release}}$. An external agent must acquire mastership of the system interface before it is allowed to issue an external request by asserting $\overline{\text{ExtRqst}}$ and waiting for the processor to assert $\overline{\text{Release}}$ for one cycle. The processor will not assert $\overline{\text{Release}}$ until it is ready to accept an external request. Mastership of the system interface is always returned to the processor after an external request has been issued, and the processor will not accept a subsequent external request until it has finished the current one.

Processor requests are managed by the processor in two distinct modes, **secondary cache mode** and **no-secondary cache mode**, which reflect the presence or absence of a secondary cache, programmable via the boot time mode control interface. The allowed modes of operation are dependent on the package configuration for the processor. A processor in the large configuration package may be programmed to run in secondary cache mode or no-secondary cache mode.

In no-secondary cache mode, the processor will issue requests in a strict sequential fashion; that is, the processor is only allowed to have one request pending at any time. The processor will issue a read request and wait for the read response before issuing any subsequent requests. The processor will issue a write request only if there are no reads pending.

The processor provides the signals $\overline{\text{RdRdy}}$ and $\overline{\text{WrRdy}}$ to allow an external agent to manage the flow of processor requests. $\overline{\text{RdRdy}}$ controls the flow of processor read, invalidate, and update requests while $\overline{\text{WrRdy}}$ controls the flow of processor write requests. Processor null write requests must always be accepted, they cannot be delayed by either $\overline{\text{RdRdy}}$ or $\overline{\text{WrRdy}}$. The processor samples the signal $\overline{\text{RdRdy}}$ to determine the issue cycle (cycle in which the address cycle is accomplished and the data cycle can be issued next) for a processor read, invalidate, or update request

and the processor samples the signal $\overline{\text{WrRdy}}$ to determine the issue cycle for a processor write request. The issue cycle for a processor read, invalidate or update request is defined to be the first address cycle for the request for which the signal $\overline{\text{RdRdy}}$ was asserted two cycles previously. The issue cycle for a processor write request is defined to be the first address cycle for the write request for which the signal $\overline{\text{WrRdy}}$ was asserted two cycles previously. If the processor wishes to issue a request but is unable to because one of the signals $\overline{\text{RdRdy}}$ or $\overline{\text{WrRdy}}$ is de-asserted, the processor will repeat the address cycle for the request until the issue cycle is accomplished. Once the issue cycle is accomplished, data transmission will begin for a request that includes data. There will always be one and only one issue cycle for any processor request.

The processor will accept external requests while attempting to issue a processor request by releasing the system interface to slave state in response to an assertion of $\overline{\text{ExtRqst}}$. Note that the rules governing the issue cycle of a processor request are strictly applied to determine the action the processor is taking. The processor will either accomplish the issue of the processor request, in which case the processor request will be completed in its entirety before an external request will be accepted, or the processor will release the system interface to slave state without accomplishing the issue of the processor request. In the latter case, the processor will attempt to issue the processor request again after the external request is completed, and the rules governing issue cycle will again apply.

In the no-secondary cache mode an external agent must be capable of accepting a processor read request at any time there are no processor read requests pending and the signal $\overline{\text{RdRdy}}$ has been asserted for two or more cycles. An external agent must be capable of accepting a processor write request at any time there are no processor read requests pending and the signal $\overline{\text{WrRdy}}$ has been asserted for two or more cycles.

In the secondary cache mode, the processor will issue requests both individually as in no-secondary cache mode and in groups that begin with processor read requests called clusters. Specifically, the processor will issue individual read requests, invalidate or update requests, and write requests and the processor will issue clusters. A cluster consists of a processor read request followed by one or two additional processor requests issued while the read request is pending. All of the requests that are part of a cluster must be accepted before the response to the read request that begins the cluster may be returned to the processor. A cluster will consist of a processor read request followed by a write request, a processor read request followed by a potential update request, or a processor read request followed by a potential update request, followed by a write request.

The issue of potential update requests within a cluster can be disabled via the boot time mode control interface. A processor potential update request is defined to be any update request that is issued while a processor read request is pending. In addition, a bit in the command for processor updates identifies potential updates. Potential updates are issued in conjunction with a processor read request. That is, once the processor accomplishes the issue of a read request, a potential update request will follow if one is required regardless of the state of $\overline{\text{RdRdy}}$. Potential update requests do not obey the $\overline{\text{RdRdy}}$ flow control rules for issue, but rather issue with a single address cycle regardless of the state of $\overline{\text{RdRdy}}$.

A write request that is part of a cluster does obey the $\overline{\text{WrRdy}}$ rules for issue, and the processor will accept external requests between the issue of a processor read request, or a processor read request followed by a potential update request, and the issue of a processor write request within a cluster. The processor signals that it is issuing a cluster that contains a processor write request by issuing a read with write forthcoming request instead of an ordinary read request to start the cluster. The read with write forthcoming request is identified by a bit in the command for processor read requests. The external agent must accept all of the requests that form a cluster before it may return a response to the read request that began the cluster. The behavior of the processor is undefined if the external agent returns a response to a processor read request that begins a cluster before accepting all of the requests that form the cluster.

Since the processor will accept external requests between the issue of a read with write forthcoming request that begins a cluster and the issue of the write request that completes the cluster, it is possible for an external request to obviate the need for the write request within the cluster. For instance, if the external agent were to issue an external invalidate request that targeted the cache line the processor was attempting to write back, the state of the cache line would be changed to invalid, and the write back for the cache line would no longer be needed. In this event, the processor will issue a processor null write request after completing the external request to complete the cluster.

Processor null write requests do not obey the $\overline{\text{WrRdy}}$ flow control rules for issue, but rather issue with a single address cycle regardless of the state of $\overline{\text{WrRdy}}$. Any external request that changes the state of a cache line from dirty exclusive or dirty shared to clean exclusive, shared, or invalid will obviate the need for a write back of that cache line.

A processor potential update request remains potential until the response to the pending processor read request that began the cluster is received. If the read response data is returned in one of the shared states, shared or dirty shared, the potential update is no longer potential and must receive an acknowledge via either the signal lvdAck or lvdErr . If the read response data is returned in one of the exclusive states, clean exclusive or dirty exclusive, the potential update is nullified and the processor will neither expect nor require an acknowledge.

In secondary cache mode, an external agent must be capable of accepting a processor read request followed by a potential update request at any time there are no processor read requests pending, no unacknowledged processor invalidate or update requests, and the signal $\overline{\text{RdRdy}}$ has been asserted for two or more cycles. An external agent must be capable of accepting a processor write request at any time there are no processor read requests pending, there is a processor read with write forthcoming request pending, no unacknowledged processor invalidate or update requests that are compulsory, and the signal $\overline{\text{WrRdy}}$ has been asserted for two or more cycles.

After issuing a processor read request, the processor will not attempt to issue a subsequent read request until it has received a read response for the read request, whether it began a cluster or not. After issuing a processor invalidate or update request, or after a potential update request becomes compulsory, the processor will not attempt to issue a subsequent request until it has received an acknowledge for the invalidate or update request. After the processor has issued a write request, the processor will not attempt to issue a subsequent request until at least four cycles after the issue cycle of the write request.

The following sections detail the sequence, protocol and syntax of processor and external requests. Sequence refers to the precise series of requests that a processor generates to service a system event. Protocol refers to the cycle by cycle signal transitions that occur on the processor's system interface pins to realize a processor or external request. Syntax refers to the precise definition of bit patterns on encoded buses such as the command bus.

4.2 Processor Request Sequencing

The processor will generate a request or a series of requests through the system interface to satisfy system events. Processor requests are managed in two distinct modes, secondary cache mode and no-secondary cache mode. The following sections detail the sequence of requests generated by the processor for each system event in secondary cache and no-secondary cache mode.

4.2.1 Primary and Secondary Cache Miss on a Load

When the processor misses in both the primary and secondary caches on a load, it must obtain the cache line that contains the data element to be loaded from an external agent before it can proceed. If the new cache line will replace a current cache line that is in the state dirty exclusive or dirty shared, the current cache line must be written back before the new line can be loaded in the primary and secondary caches.

The processor will examine the coherency attribute in the TLB entry for the page that contains the requested cache line and if the coherency attribute is exclusive it will issue a coherent read request that also requests exclusivity.

If the coherency attribute is sharable or update the processor will issue a coherent read request and if the coherency attribute is noncoherent the processor will issue a noncoherent read request.

In no-secondary cache mode, the processor will issue a read request for the cache line that contains the data element to be loaded, wait for an external agent to provide the read data in response to the read request, and then, if the current cache line must be written back, the processor will issue a write request for the current cache line.

In secondary cache mode, the processor will issue a read request for the cache line that contains the data element to be loaded if the current cache line does not need to be written back and the coherency attribute for the page that contains the requested cache line is anything other than exclusive. If the current cache line needs to be written back and the coherency attribute for the requested cache line is not exclusive, the processor will issue a cluster consisting

of a read with write forthcoming request for the cache line that contains the data element to be loaded followed by a write request for the current cache line. If the current cache line needs to be written back and the coherency attribute for the page that contains the requested cache line is exclusive, the processor will issue a cluster consisting of an exclusive read with write forthcoming request followed by a write request for the current cache line. If the current cache line needs to be written back and the coherency attribute for the page that contains the requested cache line is update on write, the processor will issue a cluster consisting of a read with write forthcoming request followed by a write request for the current cache line.

The processor request for a primary and secondary cache miss on a load is tabulated below.

Table 4-1 Primary and Secondary Cache Miss on a Load

Page Attribute	Processor Configuration	No Secondary Cache Mode		Secondary Cache Mode	
		State of the Data Cache Line Being Replaced			
		Clean/Invalid	Dirty	Clean/Invalid	Dirty
Noncoherent	All VR4000 and VR4400	NCR	NCR/W	NCR	NCR-W
Exclusive (Write Invalidate)	VR4000SC, VR4400SC, VR4400MC	NA	NA	R _{Ex}	R _{Ex} -W
Sharable (Write Invalidate)	VR4400MC	NA	NA	R	R-W
Update (Write Update)	VR4400MC	NA	NA	R	R-W

Remark	NCR	Processor Non-Coherent Read Request
	NCR/W	Processor Non-Coherent Read Request followed by Processor Write Request
	NCR-W	Cluster. Processor Non-Coherent Read Request with Write Forthcoming followed by Processor Write Request
	R	Processor Coherent Read Request
	R-W	Cluster. Processor Coherent Read Request with Write Forthcoming followed by Processor Write Request
	REx	Processor Read Request with Exclusivity
	REx-W	Cluster. Processor Read Request with Exclusivity and Write Forthcoming followed by Processor Write Request
	NA	Not Available

4.2.2 Primary and Secondary Cache Miss on a Store

When the processor misses in both the primary and secondary caches on a store, it must obtain the cache line that contains the target location of the store from an external agent before it can proceed. In secondary cache mode, if the new cache line will replace a current cache line that is in the state dirty exclusive or dirty shared, the current cache line must be written back before the new line can be loaded in the primary and secondary caches. In non-secondary cache mode, if the new cache line will replace a current cache line that is in the state dirty exclusive, the current cache will be moved to an internal write buffer before the new cache line is loaded in the primary cache.

The processor will examine the coherency attribute in the TLB entry for the page that contains the requested cache line to see if this cache line is being maintained with a write invalidate or a write update cache coherency protocol. If the coherency attribute is sharable or exclusive a write invalidate protocol is in effect, and a coherent read that also requests exclusivity will be issued. If the coherency attribute is update a write update protocol is in effect and a coherent read request will be issued. If the coherency attribute is noncoherent a noncoherent read request will be issued.

In no-secondary cache mode, the processor will issue a read request for the cache line that contains the data element to be loaded, wait for an external agent to provide the read data in response to the read request, and then, if the current cache line must be written back, the processor will issue a write request for the current cache line.

In secondary cache mode, the processor will issue a read request for the cache line that contains the target location of the store if the current cache line does not need to be written back and the coherency attribute for the page that contains the requested cache line is noncoherent. If the current cache line does not need to be written back and the coherency attribute for the page that contains the requested cache line is sharable or exclusive, the processor will issue a cluster consisting of a read request. If the current cache line does not need to be written back and the coherency attribute for the page that contains the requested cache line is update, and potential updates are enabled, the processor will issue a cluster consisting of a read request followed by a potential update request. If the current cache line needs to be written back and the coherency attribute for the requested cache line is noncoherent, the processor will issue a cluster consisting of a read with write forthcoming request for the cache line that contains the target location of the store followed by a write request for the current cache line. If the current cache line needs to be written back and the coherency attribute for the page that contains the requested cache line is sharable or exclusive, the processor will issue a cluster consisting of a read with write forthcoming request, followed by a write request for the current cache line. If the current cache line needs to be written back and the coherency attribute for the page that contains the requested cache line is update and potential updates are enabled, the processor will issue a cluster consisting of a read with write forthcoming request, followed by a potential update, followed by a write request for the current cache line.

If the processor issues a cluster that contains a potential update, and the response data for the read request is returned with an indication that it must be placed in the cache in a shared state, either shared or dirty shared, the potential update becomes compulsory. Once a potential update becomes compulsory, the external agent must forward the update to the system, and signal an acknowledge to the processor when the update is complete. In this case the processor will not complete the store until the update has been acknowledged.

If the processor issues a cluster that contains a potential update, and the response data for the read request is returned in an exclusive state, clean exclusive or dirty exclusive, the potential update is nullified. Once a potential update has been nullified, the external agent must simply discard the update. The processor will not wait for or expect an acknowledge to a potential update that has been nullified.

If the processor issues a read request, or a cluster that does not contain a potential update, and the response data for the read request is returned with an indication that it must be placed in the cache in a shared state, either shared or dirty shared, the processor will then issue an invalidate request or an update request depending on the coherency attribute for the page that contains the target location of the store instruction. If the coherency attribute is update, an update request will be issued, otherwise an invalidate request will be issued. The external agent must forward the invalidate or update to the system and signal an acknowledge to the processor for the invalidate or update request. The processor will not complete the store until it has received an acknowledge for the invalidate or update request.

The concept of potential updates is introduced to provide the external agent a chance to use the system bus more efficiently. In an update protocol, it is quite likely that a cache line requested by a processor coherent read request will be returned in a shared state, and that the processor will then have to issue an update request before it can complete a store instruction. The potential update issued with the read request in a cluster allows the external agent to anticipate the read response on the system bus, and if it arrives with an indication that it is shared to quickly gain control of the system bus and transmit the required update to the rest of the system. This provides the soonest possible acknowledge to the processor and allows the processor to complete the store instruction as quickly as possible. Without the potential update request the response data must be returned to the processor, after which the processor will issue an update request which must then be forwarded to the system bus before an acknowledge can be returned to the processor.

Note that potential updates behave in all cases as if they have not yet been issued by the processor. Potential updates are not subject to cancellation, and do not expect or require an acknowledge. When a potential update is nullified, the processor behaves as if no update request was ever issued. When a potential update becomes compulsory, the processor behaves as if it had issued an update request at that instant. Once a potential update becomes compulsory it is subject to cancellation, and the processor requires an acknowledgment.

The processor request for a primary and secondary cache miss on a store is tabulated below.

Table 4-2 Primary and Secondary Cache Miss on a Store

Page Attribute	Processor Configuration	No Secondary Cache Mode		Secondary Cache Mode			
		State of the Data Cache Line Being Replaced					
		Clean/Invalid	Dirty	Clean/Invalid		Dirty	
Noncoherent	All Vr4000 and Vr4400	NCR	NCR/W	NCR		NCR-W	
Exclusive (Write Invalidate)	Vr4000SC, Vr4400SC, Vr4400MC	NA	NA	REx		REx-W	
Sharable (Write Invalidate)	Vr4400MC	NA	NA	REx		REx-W	
Update (Write Update)	Vr4400MC	NA	NA	Dis R/U	En R-PU	Dis R-W/U	En R-PU-W

Remark	NCR	Processor Non-Coherent Read Request
	NCR/W	Processor Non-Coherent Read Request followed by Processor Write Request
	NCR-W	Cluster. Processor Non-Coherent Read Request with Write Forthcoming followed by Processor Write Request
	REx	Processor Coherent Read Request with Exclusivity
	REx-W	Cluster. Processor Coherent Read Request with Exclusivity and Write Forthcoming followed by Processor Write Request
	R/U	Processor Coherent Read Request followed by Processor Update Request
	R-PU	Cluster. Processor Coherent Read Request followed by Processor Potential Update Request
	R-PU-W	Cluster. Processor Coherent Read Request followed by Processor Potential Update Request followed by Processor Write Request
	R-W/U	Cluster. Processor Coherent Read Request with Write Forthcoming followed by Processor Write Request, followed by a Processor Update Request if the read response data is shared.
	Dis	Potential Update Disable (Modebit [20]: PotUpdDis = 1)
	En	Potential Update Enable (Modebit [20]: PotUpdDis = 0)
	NA	Not Available

4.2.3 Secondary Cache Hit on a Store to a Shared Line

When the processor store hits in the secondary cache on a line that is marked shared or dirty shared, an invalidate or update request must be issued and the processor must receive an acknowledge before the store can be completed. The processor will check the coherency attribute in the TLB for the page that contains the cache line that is the target of the store to determine if the cache line is being managed using a write invalidate or write update cache coherency protocol. If the coherency attribute is sharable or exclusive a write invalidate protocol is in effect, and the processor will issue an invalidate request. If the coherency attribute is update a write update protocol is in effect, and the processor will issue an update request. The processor will not complete the store until an external agent signals an acknowledge for the invalidate or update request.

4.2.4 Uncached Load or Store

When the processor performs an uncached load, it will issue a noncoherent read request. When the processor performs an uncached store, it will issue a write request.

4.2.5 Uncached Store Buffer

On the V4400MC, there is a single entry uncached store buffer. This means when an uncached store reaches the WB stage of the pipeline, the store drops into the uncached store buffer and the CPU pipeline may continue to run. If another uncached store reaches the WB stage of the pipeline before the first uncached store has been sent out of the chip, the CPU will stall until the uncached store buffer has completed the first uncached store.

If other useful instructions can be scheduled between successive uncached stores, then considerable performance improvement may be achieved using this mechanism.

To avoid CPU stalls, for sequential stores, there needs to be 7 cycles between uncached stores:

```

                SW      r2, (r3) # uncached store
1:             NOP
2:             NOP
3:             NOP
4:             NOP
5:             NOP
6:             NOP
7:             NOP
                SW      r2, (r3) # uncached store

```

Remark NOPs can be replaced by instructions with no uncached access.

Remember that if any of the instructions between the stores are multiple cycles then the number of instructions between the stores can be reduced further without causing stalls (for example, variable shifts take 2 PCycles).

If a sequence of uncached stores are executed in a loop, then the 2 killed cycles that are lost in branch latency still count towards the 7 cycles needed between the stores:

```

loop:      SW      r2, (r3) # uncached store
1:         NOP
2:         NOP
3:         NOP
4:         BR      LOOP
5:         NOP
6:         killed
7:         killed

```

In this case only 4 instructions have to be scheduled (excluding the branch).

The latency between stores is really defined by the fact that back to back uncached stores on the system interface can be sent out at a maximum rate of one store every 4 external cycles (Address, Data, X, X). This translates to 8 pipeline cycles for a chip running the system interface in divide by 2 mode. The above cycle count and program examples are assumed to be in divide by 2 mode. If a larger clock divisor is used (div3, div4), then the number of cycles between successive uncached stores goes up by the same ratio (div3: 11 pipeline cycles, div4: 15 pipeline cycles) if you want to avoid CPU pipeline stalls.

The SYNC instruction is used by the V_R4400MC to ensure that any uncached store in the uncached store buffer is sent out of the chip before any load or store after the SYNC completes.

As external requests have higher priority than pending uncached stores, it is possible for an external agent to see the V_R4400MC complete cached and uncached stores out of program execution order. For example, given the sequence:

```

SW      r2, (r3) # uncached store
SW      r4, (r5) # cached store

```

Given an external intervention or snoop is sent into the chip when the uncached store has entered the uncached store buffer, but has not yet been sent out of the chip, and the cached store has run successfully through the TC stage (i.e., it has hit in the primary cache). The external agent will see the state of the internal caches after the cached store but before the result of the uncached store is seen outside the chip. The SYNC instruction may be used to avoid this case if necessary, as shown below:

```

SW      r2, (r3) # uncached store
SYNC
SW      r4, (r5) # cached store

```

4.2.6 Cache Instructions

The V_R4400MC processor provides a variety of cache instructions for use in maintaining the state and contents of the primary and secondary caches. During the execution of cache instructions the processor may issue write requests, or invalidate requests. For further details on cache instructions see V_R4000, V_R4400 USER'S MANUAL—ARCHITECTURE.

4.3 External Request Handling

An external agent must arbitrate with the processor for access to the system interface before it can issue an external request. The external agent will signal that it wishes to begin an external request and wait for the processor to signal that it is ready to accept the request before issuing any new external request. The processor will decide based on its internal state and the current state of the system interface when to accept a new external request. The processor will signal that it is ready to accept an external request based on the following criteria.

- (1) If there are no processor requests pending, the processor will decide based on its internal state whether to accept the external request, or rather to issue a new processor request. The processor may issue a new processor request while the external agent is requesting access to the system interface to issue an external request.
- (2) The processor will accept an external request after completing a processor request or a processor request cluster that is in progress.
- (3) While waiting for the assertion of $\overline{\text{RdRdy}}$ to issue a processor read request the processor will accept an external request provided that the request is delivered to the processor one or more cycles before $\overline{\text{RdRdy}}$ is asserted.
- (4) While waiting for the assertion of $\overline{\text{WrRdy}}$ to issue a processor write request the processor will accept an external request provided that the request is delivered to the processor one or more cycles before $\overline{\text{WrRdy}}$ is asserted.
- (5) While waiting for the response to a read request and after the V_R4400MC has made an uncompelled change to slave state, an external agent may submit an external request before providing the read response data.

4.4 Invalidate and Update Cancellation

An external agent may discover that a processor request for an invalidate or update cannot be completed based on state changes in the external system that have not yet been reflected into the processor's caches. An example of this in a bus based system is the case in which a processor issues an invalidate, but before the external bus interface can transmit the invalidate an invalidate is received from another processor that targets the same cache line. In this case, the processor's cache does not reflect the current state of the system, and the unacknowledged invalidate cannot be transmitted. When this occurs, the external agent must cancel the invalidate or update. The processor, upon receiving a cancellation, will process any external requests that the external agent wishes to issue and then re-examine the state of the cache to determine what action to take. In the above example, this would cause the processor to process an external request to invalidate the cache line that was the target of the store, after which the processor would re-examine the state of the cache and discover that the cache line that was the target of the store is now invalid. The processor would then process the store as a store miss and issue a read request instead of an invalidate request.

Potential updates may not be canceled until they become compulsory. Potential updates are issued within a cluster under pending reads and are no longer potential after the read request is satisfied. In more general terms, an external request that indicates processor update cancellation may not be issued when a processor read is pending and may not be issued unless a compulsory update is unacknowledged. The behavior of the processor is undefined if the cancellation indication is signaled on an external coherence request to the processor while a processor read is pending or when there is no compulsory update unacknowledged.

4.5 Load Linked Store Conditional Considerations

Generally the execution of a load linked store conditional instruction sequence is not visible at the system interface, that is no special requests are generated due to the execution of this instruction sequence. However, there is one situation for which the execution of a load linked store conditional instruction sequence will be visible as a change in the nature of a processor read request.

Specifically, if the data location targeted by a load linked store conditional instruction sequence maps to the same cache line that the instruction area containing the load linked store conditional code sequence is mapped to, then immediately after executing the load linked instruction the cache line that contains the link location will be replaced by the instruction line containing the code. The link address is kept in a register separate from the cache and remains active as long as the link bit remains set. The link bit is set by the load linked instruction, and is cleared by any change of cache state for the cache line containing the link address, or a return from exception.

In order for the load linked store conditional instruction sequence to work correctly all coherency traffic targeting the link address must be visible to the processor, and the cache line containing the link location must remain in a shared state in every cache in the system. This guarantees that a store conditional executed by some other processor is visible to the processor as a coherence request which changes the state of the cache line that contains the link location. To accomplish this, a read request issued by the processor which causes the replacement of a cache line that contains the link location while the link bit is set will indicate that the link address is being retained. The link address retained bit in the command for the read request will be asserted to provide this indication. This informs the external agent that even though the processor has replaced this cache line and no longer has it present in its cache, it still must see any coherence traffic that targets this cache line.

In addition, any snoop or intervention request that targets a cache line which is not present in the cache, but for which the snoop or intervention address matches the current link address while the link bit is set, will return an indication that the cache line is present in the cache in a shared state. A shared indication is returned even though the processor does not actually have the data content of the cache line. This is consistent since the processor never returns data in response to an intervention request for a cache line that is in the shared state. The shared response guarantees that the cache line that contains the link location will remain in a shared state in all other processor's caches, and therefore that any other processor that attempts a store conditional to this link location must issue a coherence request in order to complete the store conditional.

The address value of the LLAddr register is lost in the following cases:

- If the lines including link addresses are invalidated by an external request (external invalidate, snoop, or intervention request).
- When ERET instruction is completed
- When the cache lines including link addresses are updated by an external update request

4.6 System Interface Endianness

The endianness of the system interface is programmed at boot time via the boot time mode control interface and is fixed until the next time the processor mode bits are read. The endianness of the system interface and the external system cannot be changed by software. The reverse endian bit can be set by software to reverse the interpretation of endianness inside the processor, but the endianness of the system interface remains unchanged. For further details on the reverse endian bit see V_R4000, V_R4400 USER'S MANUAL—ARCHITECTURE.

4.7 System Interface Protocol

The system interface protocol describes the cycle by cycle signal transitions that occur on the pins of the system interface to realize requests between the processor and an external agent.

4.7.1 Introduction

The system interface is register to register. That is, processor outputs come directly from output registers and begin to change with the rising edge of SClock and processor inputs are fed directly to input registers that latch the inputs with the rising edge of SClock. Therefore, if an input to the processor is changed during a particular cycle such that the new value is sampled at the end of the cycle, the earliest the processor can change one of its outputs in response to the input change is two cycles later. This methodology was chosen to allow the system interface to run at the highest possible clock frequency.

The primary communication paths for the system interface are a sixty-four bit address and data bus, SysAD(63:0) and a nine bit command bus, SysCmd(8:0). The SysAD bus and the SysCmd bus are bidirectional, that is they are driven by the processor to issue a processor request, and by an external agent to issue an external request. When the processor is driving the SysAD bus and the SysCmd bus the system interface is in **master state**, when an external agent is driving the SysAD bus and the SysCmd bus the system interface is in **slave state**.

A request through the system interface consists of an address, a system interface command that specifies the precise nature of the request, and a series of data elements if the request is for a write, read response, or update. Addresses and data elements are transmitted on the SysAD bus. System interface commands and data identifier are transmitted on the SysCmd bus.

Cycles in which the SysAD bus contains a valid address are called address cycles and cycles in which the SysAD bus contains a valid data element are called data cycles. In master state the processor will assert the signal ValidOut whenever the SysAD bus and the SysCmd bus are valid. In slave state an external agent will assert the signal ValidIn whenever the SysAD bus and the SysCmd bus are valid.

The SysCmd bus is used to identify the contents of the SysAD bus during any cycle in which it is valid. The most significant bit, SysCmd(8), is always used to indicate whether the current cycle is an address cycle or a data cycle. During address cycles, SysCmd(7:0) will contain a system interface command. The encoding of system interface commands is detailed in the section on system interface syntax. During data cycles, SysCmd(7:0) will contain an indication of whether this is the last data element to be transmitted and other information about the data element. The contents of the SysCmd bus during data cycles is called a **data identifier**. The encoding of data identifiers is detailed in 4.9 System Interface Syntax.

A request through the system interface consists of one or more identical address cycles, followed by a series of data cycles for requests that include data. The most efficient request through the system interface will consist of a single address cycle followed by a single data cycle or a number of data cycles sufficient to transmit a block of data.

4.7.2 System Interface Arbitration

When an external agent needs to issue an external request through the system interface, it must first get the system interface into slave state. The transition from master state to slave state is arbitrated by the processor using the system interface handshake signals ExtRqst and Release. An external agent will signal that it wishes to issue an external request by asserting ExtRqst, and the processor will release the system interface from master state to slave state by asserting Release for one cycle when it is ready to accept an external request. The system interface will return to master state as soon as the issue of the external request is complete. Having asserted ExtRqst, an external agent must not de-assert ExtRqst until the processor asserts Release. After the processor asserts Release, the external agent should de-assert ExtRqst no more than two cycles after the assertion of Release. An external agent may continue to assert ExtRqst if another external request follows the current request. After the first external request completes, the processor must assert Release again before the second external request is issued to the processor.

The system interface will remain in master state until the external agent requests and is granted the system interface or until the processor issues a read request, or completes the issue of a cluster. Whenever a processor read request is pending, after the issue of a read request or after the issue of all of the requests in a cluster, the processor will switch the system interface to slave state even though the external agent is not arbitrating to issue an external request. This transition to slave state is specifically to allow the external agent to return read response data. The external agent must not assert the signal ExtRqst for the purposes of transitioning the system interface to slave state to return read response data. ExtRqst will only be asserted when the external agent needs to get the system interface into slave state so that it can issue an external request.

The signal ExtRqst is strictly used to arbitrate for the system interface, that is to request the transition of the system interface from master state to slave state. ExtRqst should be de-asserted no more than two cycles after a cycle in which Release is asserted unless the external agent wishes to perform a subsequent external request.

The transition of the system interface from master state to slave state initiated by the processor when a processor read request is pending will be referred to as an uncompelled change to slave state. An uncompelled change to slave state will occur during or some number of cycles after the issue cycle of a read request or the last cycle of the last request in a cluster. The number of cycles depends on the state of the cache, the presence of a secondary cache and the secondary cache parameters. After an uncompelled change to slave state the system interface will remain in slave state until the external agent issues an external request, after which the system interface will return to master state. An external agent must note that the processor has performed an uncompelled change to slave state and begin driving the address and data bus and the command bus. As long as the system interface is in slave state, the external agent will begin an external request without arbitrating for the system interface, that is without asserting ExtRqst.

4.7.3 System Interface Signal Descriptions

The system interface address and data bus is the SysAD bus. The system interface command bus is the SysCmd bus.

The SysAD bus and SysCmd bus valid signal that is asserted by the processor in master state is ValidOut. The SysAD bus and SysCmd bus valid signal that is asserted by an external agent in slave state is ValidIn.

The SysADC bus provides eight check bits for the SysAD bus. The nature of the check bits is programmable via the boot time mode control interface. The check bits may represent even byte parity for the contents of the SysAD bus, or they may be interpreted according to the code described in CHAPTER 5 ERROR CHECKING AND CORRECTING (ECC) to detect and correct single bit errors and detect double bit errors or three or four bit errors within a nibble on the SysAD bus. For a description of even parity, see APPENDIX B EVEN PARITY. For further details on the ECC characteristics of the Vn4400MC, see CHAPTER 5 ERROR CHECKING AND CORRECTING (ECC).

The signal SysCmdP is an even parity bit over the nine bits of the SysCmd bus generated by the processor in master state. SysCmdP is not checked by the processor in slave state. For a description of even parity, see APPENDIX B EVEN PARITY.

System interface arbitration is implemented using the signals ExtRqst and Release.

Processor request flow control is implemented using the signals RdRdy and WrRdy.

Processor invalidate and update requests are acknowledged using IvdAck and IvdErr signals.

4.7.4 System Interface Maintenance Signals

In addition to the signals used to implement the system interface request protocol, the system interface includes maintenance signals necessary for the operation of the processor. These include a master clock input for the processor, MasterClock, which must be connected to a continuous clock signal at the desired operation frequency of the processor; a processor synchronization clock output, SyncOut and a processor synchronization clock input, SyncIn that must be connected together to allow the processor internal clock synchronization logic to compensate for pad driver and receiver delays; a master clock output, MasterOut, aligned with MasterClock for use in clocking external logic that must run at MasterClock frequency; three reset related inputs, VppOk, ColdReset, and Reset; an

eight bit status bus, **Status(7:0)** that is encoded to indicate the current operation status of the processor; a system fault processor output, **Fault** that is the mismatch indication of the boundary comparators when the processor is running in master checker mode; two transmit clock outputs, **TClock(1:0)** and two receive clock outputs, **RClock(1:0)** at the programmed operation frequency of the system interface.

4.7.5 System Interface Signal Summary

SysAD(63:0):	(i/o)	A 64-bit bus used for address and data transmission between the processor and an external agent.
SysADC(7:0):	(i/o)	An 8-bit bus containing check bits for the SysAD bus.
SysCmd(8:0):	(i/o)	A 9-bit bus used for command and data identifier transmission between the processor and an external agent.
SysCmdP:	(i/o)	A single even parity bit over the SysCmd bus.
ValidIn:	(i)	Signals that an external agent is driving a valid address or valid data on the SysAD bus and a valid command or data identifier on the SysCmd bus during this cycle.
ValidOut:	(o)	Signals that the processor is driving a valid address or valid data on the SysAD bus and a valid command or data identifier on the SysCmd bus during this cycle.
ExtRqst:	(i)	Signals that an external agent wishes to issue an external request.
Release:	(o)	Signals that the processor is releasing the system interface to slave state.
RdRdy:	(i)	Signals that an external agent is capable of accepting a processor read, invalidate, or update request in both no-secondary cache and secondary cache mode or a read followed by a potential update request in secondary cache mode.
WrRdy:	(i)	Signals that an external agent is capable of accepting a processor write request in both no-secondary cache and secondary cache mode.
IvdAck:	(i)	Signals that a processor invalidate or update request has completed successfully.
IvdErr:	(i)	Signals that a processor invalidate or update request has completed unsuccessfully.
TClock(1:0):	(o)	Two identical transmit clocks at the operation frequency of the system interface.
RClock(1:0):	(o)	Two identical receive clocks at the operation frequency of the system interface.
MasterClock:	(i)	Master clock input at the operation frequency of the processor.
MasterOut:	(o)	Master clock output aligned with MasterClock.
SyncOut:	(o)	Synchronization clock output.
SyncIn:	(i)	Synchronization clock input.
Status(7:0):	(o)	An 8-bit bus that indicates the current operation status of the processor.
VddOk:	(i)	When asserted, this signal indicates to the V ₄₄₀₀ MC that the power supply voltage has been within the specified range for more than 100 milliseconds and will remain stable. The assertion of VddOk will initiate the reading of the boot time mode control serial stream.
ColdReset:	(i)	This signal must be asserted for a power on reset or a cold reset. The clocks SClock, TClock, and RClock begin to cycle and are synchronized with the deassertion edge of ColdReset.
Reset:	(i)	This signal must be asserted for any reset sequence. It may be asserted synchronously or asynchronously for a cold reset, or synchronously to initiate a warm reset.
Fault:	(o)	Mismatch output of boundary comparators.

4.7.6 System Interface Request Descriptions

The following sections will illustrate the protocol of each processor and external request through text and detailed timing diagrams. The timing diagrams use abbreviations to show the contents of encoded busses during cycles in which they are defined. Following is a list of abbreviations used for each bus and their definitions.

Global:

Unsd – Unused.

SysAD bus:

Addr – Physical address.

Data<n> – Data element number n of a block of data.

SysCmd bus:

Cmd – An unspecified system interface command.

Read – A read request command.

RwWF – A read with write forthcoming request command.

Write – A write request command.

Null – A null request command.

SINull – A system interface release null request command.

SCNull – A secondary cache release null request command.

Ivd – An invalidate request command.

Upd – An update request command.

Ivtn – An intervention request command.

Snoop – A snoop request command.

NData – A noncoherent data identifier for a data element other than the last data element.

NEOD – A noncoherent data identifier for the last data element.

CData – A coherent data identifier for a data element other than the last data element.

CEOD – A coherent data identifier for the last data element.

4

Two closely spaced wavy vertical lines in a timing diagram indicate a repetition of the current cycle. That is the cycle broken by the wavy lines may represent one or more identical cycles. This is referred to as a break in the timing diagram and is used to keep the timing diagrams concise and readable.

4.7.7 Invalidate and Update Acknowledge Protocol

Processor invalidate and update requests are acknowledged using the signals IvdAck and IvdErr. An external agent will drive either IvdAck or IvdErr for one cycle to signal the completion status of the current processor invalidate or update request. Invalidate or update request acknowledge occurs in parallel with requests on the SysAD and SysCmd busses. IvdAck or IvdErr may be driven at any time after a processor invalidate or update request is issued provided that the update request is compulsory.

4.7.8 Arbitration Protocol


System interface arbitration is implemented using the signals ExtRqst and Release. When an external agent wishes to issue an external request, it will assert ExtRqst. The processor will wait until it is ready to handle an external request and assert Release for one cycle before it 3-states the SysAD bus and SysCmd bus. The external agent will begin driving the SysAD bus and the SysCmd bus two cycles after a cycle in which Release is asserted. The external agent should always deassert ExtRqst no more than two cycles after a cycle in which Release is asserted unless the external agent wishes to perform a subsequent external request. The external agent will always release the SysAD bus and the SysCmd bus at the completion of an external request.

The processor will assert Release for one cycle as a processor read request is issued or sometimes after a processor read request is issued to perform an uncompelled change to slave state. An external agent must begin

driving the SysAD bus and the SysCmd bus at least two cycles after the cycle in which Release is asserted. After an uncompelled change to slave state, the processor will return to master state at the end of the next external request, which may be the read response, or may be some other external request.

The processor to system handshake for external requests is illustrated in Fig. 4-1 Arbitration Protocol for External Requests.

Fig. 4-1 Arbitration Protocol for External Requests



4.7.9 Processor Read Request Protocol

A processor read request is issued with the system interface in master state by driving a read command on the SysCmd bus and a read address on the SysAD bus and asserting ValidOut for one cycle. Only one processor read request may be pending at a time. The processor must wait for and retire an external read response before initiating a subsequent read.

The processor will make an uncompelled change to slave state either at the issue cycle of the read request or sometime after the issue cycle of the read request by asserting the Release signal for one cycle. Once in slave state, an external agent may return the requested data via a read response. An external agent must not assert the signal ExtRqst for the purposes of returning a read response, but rather must wait for the uncompelled change to slave state. The signal ExtRqst may be asserted before or during a read response for the purposes of performing an external request other than a read response.

When a read is pending, ExtRqst is asserted, and Release is asserted for one cycle it may be unclear if this assertion of Release is in response to ExtRqst, or represents an uncompelled change to slave state. The only situation in which this assertion of Release may not be considered an uncompelled change to slave state is if the system interface is operating in secondary cache mode, the read request was a read with write forthcoming request, and the expected write request has not yet been issued by the processor. In this case, the write request must be accepted by the external agent before the read response can be issued. In all other cases, the assertion of Release may be considered to be an uncompelled change to slave state or to be in response to the assertion of ExtRqst. In this situation, the processor will accept either a read response, or any other external request. If an external request other than a read response is issued, the processor will perform another uncompelled change to slave state after processing of the external request is completed.

The read response may either return the requested data, or an indication that the returned data is erroneous, if the requested data could not be successfully retrieved, which will cause the processor to take a bus error.

A processor read request and an uncompelled change to slave state occurring as the read request is issued is illustrated in Fig. 4-2 Processor Read Request Protocol. A processor read request and the subsequent uncompelled change to slave state occurring sometime after the read request is issued is illustrated in Fig. 4-3 Processor Read Request Protocol, Change to Slave State Delayed.

Fig. 4-2 Processor Read Request Protocol





Fig. 4-3 Processor Read Request Protocol, Change to Slave State Delayed



4.7.10 Processor Write Request Protocol

Processor write requests are issued with one of two protocols. Double word, word, and partial word writes use a single word write request protocol. Write requests for a block of data use a block write request protocol. Processor write requests are issued with the system interface in master state.

A processor single word write request is issued by driving a write command on the SysCmd bus and a write address on the SysAD bus and asserting **ValidOut** for one cycle, followed by driving a data identifier on the SysCmd bus and data on the SysAD bus and asserting **ValidOut** for one cycle. The data identifier associated with the data cycle must contain a last data cycle indication (NEOD or CEOD).

A processor block write request is issued by driving a write command on the SysCmd bus and a write address on the SysAD bus and asserting **ValidOut** for one cycle, followed by driving a data identifier on the SysCmd bus and data on the SysAD bus and asserting **ValidOut** for a number of cycles sufficient to transmit the block of data. The data identifier associated with the last data cycle must contain a last data cycle indication. The first data cycle may not immediately follow the address cycle. A processor noncoherent single word write request is illustrated in Fig.

4-4 Processor Noncoherent Single Word Write Request Protocol. A processor block write request for eight words of data is illustrated in Fig. 4-5 Processor Coherent Block Write Request Protocol (a) and Fig. 4-6 Processor Coherent Block Write Request Protocol (b).

Fig. 4-4 Processor Noncoherent Single Word Write Request Protocol




Fig. 4-5 Processor Coherent Block Write Request Protocol (a)





Fig. 4-6 Processor Coherent Block Write Request Protocol (b)



4


4.7.11 Processor Invalidate and Update Request Protocol

A processor invalidate request or update request will use the same protocol as a single word write request except that the command associated with the address cycle will indicate that this is an invalidate or update request. The single data cycle will be unused for an invalidate.

4.7.12 Processor Null Write Request Protocol

A processor null write request is issued with the system interface in master state by driving a null command on the SysCmd bus and asserting ValidOut for one cycle. The SysAD bus is unused during the address cycle associated with a null write request. Processor null write requests cannot be flow controlled with either RdRdy or WrRdy, but rather always issue with a single address cycle. A processor null write request is illustrated in Fig. 4-7 Processor Null Write Request Protocol.

Fig. 4-7 Processor Null Write Request Protocol




4.7.13 Processor Cluster Protocol

In secondary cache mode, the processor will issue requests both individually as in no-secondary cache mode and in groups that begin with a processor read request called **clusters**. A cluster consists of a processor read request followed by one or two additional processor requests issued while the read request is pending. All of the requests that are part of a cluster must be accepted before the response to the read request that begins the cluster may be returned to the processor. A cluster will include a processor read request followed by any of a write request, a potential update request, or a potential update request followed by a write request.

The protocol of each of the requests that form a cluster is as described above. The number of unused cycles between the requests that form a cluster is specified in 4.8 Cycle Counts for System Interface Interactions. The processor will make an uncompelled change to slave state either during or following the last cycle of the last request in the cluster. A cluster consisting of a read request followed by a potential update request followed by a block write request for eight words of data with minimum spacing between the requests that form the cluster and an uncompelled change to slave state at the earliest opportunity is illustrated in Fig. 4-8 Processor Cluster Protocol.

Fig. 4-8 Processor Cluster Protocol



4.7.14 External Request Protocol

External requests may only be issued with the system interface in slave state. An external agent must assert $\overline{\text{ExtRqst}}$ to arbitrate for the system interface, and wait for the processor to release the system interface to slave state before issuing an external request. If the system interface is already in slave state, i.e. the processor has previously performed an uncompelled change to slave state, an external agent may begin an external request immediately.

After issuing an external request an external agent must return the system interface to master state. If the external agent does not have any additional external requests to perform, $\overline{\text{ExtRqst}}$ must be de-asserted two cycles after the cycle in which $\overline{\text{Release}}$ is asserted. An external agent may hold $\overline{\text{ExtRqst}}$ asserted if it needs to issue a string of external requests, but it must wait for the processor to assert $\overline{\text{Release}}$ and return the system interface to slave state before it may proceed with the next external request. For a string of external requests, the external agent must de-assert $\overline{\text{ExtRqst}}$ two cycles after the cycle in which $\overline{\text{Release}}$ is asserted for the last external request in the string. The processor will continue to handle external requests as long as $\overline{\text{ExtRqst}}$ is held asserted, however, the processor will not release the system interface to slave state for a subsequent external request until it has completed the current request. A string of external requests will not be interrupted by a processor request as long as $\overline{\text{ExtRqst}}$ is held asserted throughout the issue of the string of external requests.

(1) External Read Request Protocol


External read requests use a non-split protocol that does not allow any other request to occur at the system interface between the external read request and the read response. The protocol of an external read request encompasses the request from an external agent and the response from the processor.

An external read request consists of driving a read request command on the SysCmd bus and a read request address on the SysAD bus and asserting $\overline{\text{ValidIn}}$ for one cycle. After the address and command are sent, the external agent will release the SysCmd and SysAD busses and allow the processor to begin driving them. The processor will access the data that is the target of the read and return the data to the external agent. The processor accomplishes this by driving a data identifier on the SysCmd bus, the response data on the SysAD bus, and asserting $\overline{\text{ValidOut}}$ for one cycle. The data identifier will indicate that this is response data and contain a last data cycle indication. The processor will continue driving the SysCmd and SysAD busses after the read response is returned to transition the system interface back to master state.

External read requests are only allowed to read a word of data from the processor. The processor response to external read requests for any data element other than a word is undefined.

An external read request with the system interface initially in master state is illustrated in Fig. 4-9 External Read Request, System Interface in Master State.

Fig. 4-9 External Read Request, System Interface in Master State



Remark The V_R4400MC does not contain any resources that are readable with an external read request. The V_R4400MC will return a bus error response to any external read request.

(2) External Null Request Protocol

The V_R4400MC processor supports two kinds of external null requests. A system interface release external null request is used to return the system interface to master state after it has been released to slave state without affecting the processor. A secondary cache lease external null request is used to return ownership of the secondary cache to the processor while the system interface remains in slave state for some period of time. This is important since any time the processor releases the system interface to slave state to accept an external request, it also acquires ownership of the secondary cache for use by the external request in anticipation of handling a coherence request. When an external agent requests ownership of the system interface for the purposes of using the SysAD bus for a transfer unrelated to the processor this ownership of the secondary cache will prevent the processor from satisfying subsequent primary cache misses. The secondary cache release external request can be issued by the external agent to return ownership of the secondary cache to the processor.

External null requests require no action from the processor other than to return the system interface to master state or to regain ownership of the secondary cache.

An external null request consists of driving a null request command on the SysCmd bus and asserting $\overline{\text{ValidIn}}$ for one cycle. The SysAD bus is unused during the address cycle associated with an external null request. After the address cycle is issued the null request is complete. For a system interface release external null request the external agent will release the SysCmd and SysAD busses and allow the system interface to return to master state. For a secondary cache release external null request the system interface will remain in slave state. A secondary cache release external null request with the system interface initially in master state is illustrated in Fig. 4-10 **Secondary Cache Release External Null Request**. A system interface release external null request with the system interface initially in slave state is illustrated in Fig. 4-11 **System Interface Release External Null Request**.

Fig. 4-10 Secondary Cache Release External Null Request





Fig. 4-11 System Interface Release External Null Request



(3) External Write Request Protocol


External write requests use a protocol identical to the processor single write protocol except that the signal $\overline{\text{ValidIn}}$ is asserted instead of the signal $\overline{\text{ValidOut}}$. An external write request consists of driving a write command on the SysCmd bus and a write address on the SysAD bus and asserting $\overline{\text{ValidIn}}$ for one cycle, followed by driving a data identifier on the SysCmd bus and data on the SysAD bus and asserting $\overline{\text{ValidIn}}$ for one cycle. The data identifier associated with the data cycle must contain a last data cycle indication. After the data cycle is issued the write request is complete and the external agent will release the SysCmd and SysAD busses and allow the system interface to return to master state.

External write requests are only allowed to write a word of data to the processor. The behavior of the processor in response to an external write request for any data element other than a word is undefined.

An external write request with the system interface initially in master state is illustrated in Fig. 4-12 External Write Request.

4

Fig. 4-12 External Write Request




Remark The only writable resource in the Va4400MC is the processor interrupts.

(4) External Invalidate and Update Request Protocol

External invalidate and update requests use a protocol identical to that for external write requests. The data element provided with an update request may be a double word, word, or partial word. The single data cycle will be unused for an invalidate request. An external invalidate request following an uncompelled change to slave state is illustrated in Fig. 4-13 External Invalidate Request Following an Uncompelled Change to Slave State.

Fig. 4-13 External Invalidate Request Following an Uncompelled Change to Slave State

**(5) Read Response Protocol**

An external agent must return data to the processor in response to a processor read request by first waiting for the processor to perform an uncompelled change to slave state, and then returning the data via a single data cycle or a series of data cycles sufficient to transmit the requested data. After the last data cycle is issued the read response is complete and the external agent will release the SysCmd and SysAD busses and allow the system interface to return to master state. Note that the processor will always perform an uncompelled change to slave state at some time after issuing a read request.

The data identifier for the data cycles must indicate that this is response data, and the data identifier associated with the last data cycle must contain a last data cycle indication (NEOD or CEOD). For read responses to coherent block read requests, each data identifier must include an indication of the cache state in which to load the response data. The cache state provided with each data identifier must be the same and must be either clean exclusive, dirty exclusive, shared, or dirty shared. The behavior of the processor is undefined if the cache state provided with the data identifiers is changed during the transfer of the block of data, or if the cache state provided is invalid. The data identifier associated with a data cycle may indicate that the data transmitted during that cycle is erroneous, however, an external agent must return a block of data of the correct size regardless of erroneous data cycles. If a read response includes one or more erroneous data cycles, the processor will take a bus error. Read response data must only be delivered to the processor when a processor read request is pending; that is in response to a processor read request. The behavior of the processor is undefined if a read response is presented to it when there is no processor read pending. Further, if the processor issues a read with write forthcoming request, a processor write request or a processor null write request must be accepted before the read response may be returned. The behavior of the processor is undefined if the read response is returned before a processor write request is accepted.

A processor word read request followed by a word read response is illustrated in Fig. 4-14 **Processor Word Read Request Followed by a Word Read Response**. A read response for a processor block read with the system interface already in slave state is illustrated in Fig. 4-15 **Block Read Response, System Interface Already in Slave State**.

Fig. 4-14 Processor Word Read Request Followed by a Word Read Response





Fig. 4-15 Block Read Response, System Interface Already in Slave State



(6) External Intervention Request Protocol

External intervention requests use a protocol similar to that for external read requests except that a cache line size block of data may be returned along with an indication of the cache state for the cache line, depending on the state of the cache line and the value of the **data return** bit in the intervention request command.

The data return bit in the intervention request command may indicate return on dirty or return on exclusive. If the data return bit indicates return on dirty and the cache line that is the target of the intervention request is in the state dirty exclusive or dirty shared, the contents of the cache line will be returned in response to the intervention request. If the data return bit indicates return on exclusive and the cache line that is the target of the intervention request is in the state clean exclusive or dirty exclusive, the contents of the cache line will be returned in response to the intervention request. Otherwise, the response to the intervention request will not include the contents of the cache line but rather will simply indicate the state of the cache line that is the target of the intervention request. Note that if the cache line that is the target of the intervention request is not present in the cache at all, i.e. a tag comparison for the cache line at the target cache address fails, the cache line that is the target of the intervention request will be considered to be in the invalid state.

The processor will return an indication of the cache state in which a cache line was found but not its contents by driving a coherent data identifier that indicates the state of the cache line on the SysCmd bus, and asserting ValidOut for one cycle. The SysAD bus is unused during this data cycle. The data identifier will indicate that this is a response data cycle and will contain a last data cycle indication.

The processor will return the contents of a cache line along with an indication of the cache state in which it was found by issuing a sequence of data cycles sufficient to transmit the contents of the cache line. The data identifier transmitted with each data cycle will indicate the cache state in which the cache line was found and that this is response data. The data identifier associated with the last data cycle will contain a last data cycle indication. If the contents of a cache line is returned in response to an intervention request, it will be returned in sub-block order starting with the double word at the address supplied with the intervention request. For further details on sub-block ordering see APPENDIX A SUB-BLOCK ORDERING. Note, however, that if the intervention address targets the double word at the beginning of the block sub-block ordering is equivalent to sequential ordering. An external intervention request to a cache line found in the shared state with the system interface initially in master state is illustrated in Fig. 4-16 External Intervention Request, Shared Line, System Interface in Master State. An external intervention request to a cache line found in the dirty exclusive state with the system interface initially in slave state is illustrated in Fig. 4-17 External Intervention Request, Dirty Exclusive Line, System Interface in Slave State.

Fig. 4-16 External Intervention Request, Shared Line, System Interface in Master State





Fig. 4-17 External Intervention Request, Dirty Exclusive Line, System Interface in Slave State



4.7.15 External Snoop Request Protocol

External snoop requests use a protocol identical to that for external read requests, except that the processor will respond to a snoop request with an indication of the current cache state for the cache line that is the target of the snoop request instead of data. The processor accomplishes this by driving a coherent data identifier on the SysCmd bus, and asserting ValidOut for one cycle. The SysAD bus is unused during the snoop response. The processor will continue driving the SysCmd and SysAD busses after the snoop response is returned to transition the system interface back to master state.

Note that if the cache line that is the target of the snoop request is not present in the cache at all, i.e. a tag comparison for the cache line at the target cache address fails, the cache line that is the target of the snoop request will be considered to be in the invalid state.

An external snoop request issued with the system interface in master state is illustrated in Fig. 4-18 External Snoop Request, System Interface in Master State. An external snoop request issued with the system interface in slave state is illustrated in Fig. 4-19 External Snoop Request, System Interface in Slave State.

Fig. 4-18 External Snoop Request, System Interface in Master State





Fig. 4-19 External Snoop Request, System Interface in Slave State



4.7.16 Processor Request and Cluster Flow Control

The signal $\overline{\text{RdRdy}}$ may be used by an external agent to control the flow of a processor read, invalidate, or update request or a processor read request followed by a potential update request within a cluster. The processor samples the signal $\overline{\text{RdRdy}}$ to determine if the external agent is currently capable of accepting a read, invalidate, or update request, or a read request followed by a potential update request. The signal $\overline{\text{WrRdy}}$ controls the flow of a processor write request. The processor will not complete the issue of a read, invalidate, or update request, or a read request followed by a potential update request until it issues an address cycle for the request such that the signal $\overline{\text{RdRdy}}$ was asserted two cycles previously. The processor will not complete the issue of a write request until it issues an address cycle for the write request such that the signal $\overline{\text{WrRdy}}$ was asserted two cycles previously.

Two processor write requests in which the issue of the second is delayed for the assertion of $\overline{\text{WrRdy}}$ are illustrated in Fig. 4-20 Two Processor Write Requests, Second Write Delayed for the Assertion of $\overline{\text{WrRdy}}$. A processor cluster in which the issue of the read and a potential update request is delayed for the assertion of $\overline{\text{RdRdy}}$ is illustrated in Fig. 4-21 Processor Read Request Within a Cluster Delayed for the Assertion of $\overline{\text{RdRdy}}$. A processor cluster in which the issue of the write request is delayed for the assertion of $\overline{\text{WrRdy}}$ is illustrated in Fig. 4-22 Write Request Within a Cluster Delayed for the Assertion of $\overline{\text{WrRdy}}$. The issue of a processor write request delayed for the assertion of $\overline{\text{WrRdy}}$ and the completion of an external invalidate request is illustrated in Fig. 4-23 Processor Write Request Delayed for the Assertion of $\overline{\text{WrRdy}}$ and the Completion of an External Invalidate Request.

Fig. 4-20 Two Processor Write Requests, Second Write Delayed for the Assertion of $\overline{\text{WrRdy}}$




Fig. 4-21 Processor Read Request Within a Cluster Delayed for the Assertion of $\overline{\text{RdRdy}}$




Fig.4-22 Write Request Within a Cluster Delayed for the Assertion of $\overline{\text{WrRdy}}$





Fig. 4-23 Processor Write Request Delayed for the Assertion of $\overline{\text{WrRdy}}$ and the Completion of an External Invalidate Request



4.7.17 Data Rate Control

The system interface supports a maximum data rate of one double word per cycle. The maximum data rate the processor can support is directly related to the secondary cache access time, if the access time is too long, the processor will not be able to transmit and accept data at the maximum rate.

The rate at which data is delivered to the processor may be chosen by an external agent by driving data and asserting $\overline{\text{ValidIn}}$ every n cycles instead of every cycle. The processor will only interpret cycles during which $\overline{\text{ValidIn}}$ is asserted and the SysCmd bus contains a data identifier as valid data cycles. The processor will continue to accept data until the data word tagged as the last data word is received. An external agent may deliver data at any rate it chooses but must not deliver data to the processor faster than it is capable of accepting it.

Because the secondary cache is organized as a 128 bit RAM array, the processor will operate most efficiently if data is delivered to it in pairs of double words. It is most efficient to reduce the data rate by delivering a pair of double words to the processor, followed by some number of unused cycles, followed by another pair of double words. The pattern should be chosen to repeat at a rate determined by the secondary cache write cycle time. However, the processor will accept data in any pattern as long as the time between the transfer of any pair of odd numbered double words is greater than or equal to the write cycle time of the secondary cache. Double words in the transfer pattern are numbered beginning at zero such that the odd numbered words are the second, fourth, sixth, and so on words transferred.

The maximum processor data rate for each of the possible secondary cache write cycle times and the most efficient data pattern for each data rate is illustrated in Table 4-3 Maximum Processor Data Rates. In this and subsequent tables data patterns are specified using the letters "D" and "x", "D" indicates a data cycle and "x" indicates an unused cycle. A data pattern is specified as a sequence of letters, indicating a sequence of data and unused cycles that will be repeated to provide the appropriate data rate. For example, a data pattern specified by the sequence of letters "DDxx", to achieve a data rate of two words every four cycles, is a data pattern in which two data cycles are followed by two unused cycles followed by two data cycles and two unused cycles, and so on. A read response in which data is provided to the processor at a rate of two words every three cycles using the data pattern "DDx" is shown in Fig. 4-24 Read Response, Reduced Data Rate, System Interface in Slave State.

If data is delivered to the processor at a rate that exceeds the maximum the processor can support, based on the secondary cache write cycle time, the behavior of the processor is undefined. The secondary cache write cycle time is the sum of the parameters TWr1Dly, TWrSUP, and TWrRc described in the section on secondary cache write cycles. The rate at which the processor transmits data is programmable at boot time via the boot time mode control interface. The transmit data rate may be programmed to any of the data rates and data patterns listed in Table 4-4 Transmit Data Rates, as long as the programmed data rate does not exceed the maximum the processor can support, based on the secondary cache access time. If a transmit data rate is programmed that exceeds the maximum the processor can support, the behavior of the processor is undefined. A processor write request for which the processor transmit data rate has been programmed to 1 double word every two cycles using the data pattern "DDxx" is shown in Fig. 4-25 Processor Write Request, Transmit Data Rate Reduced.

4

Table 4-3 Maximum Processor Data Rates

<u>SCache Write Cycle Time</u>	<u>Max Data Rate</u>	<u>Best Data Pattern</u>
<= 4 PCycles	1 Double/1 Cycle	D
5-6 PCycles	2 Doubles/3 Cycles	DDx
7-8 PCycles	1 Double/2 Cycles	DDxx
9-10 PCycles	2 Doubles/5 Cycles	DDxxx
11-12 PCycles	1 Double/3 Cycles	DDxxx

Fig. 4-24 Read Response, Reduced Data Rate, System Interface in Slave State





Table 4-4 Transmit Data Rates

Data Rate	Data Pattern	Max SCache Access
1 Double/1 Cycle	D	4 PCycles
2 Doubles/3 Cycles	DDx	6 PCycles
1 Double/2 Cycles	DDxx	8 PCycles
1 Double/2 Cycles	DxDx	8 PCycles
2 Doubles/5 Cycles	DDxxx	10 PCycles
1 Double/3 Cycles	DDxxx	12 PCycles
1 Double/3 Cycles	DxxDxx	12 PCycles
1 Double/4 Cycles	DDxxxxx	16 PCycles
1 Double/4 Cycles	DxxxDxxx	16 PCycles

Fig. 4-25 Processor Write Request, Transmit Data Rate Reduced



4.7.18 Multiple Drivers on the SysAD Bus

In most V4400MC applications the SysAD bus will be a point to point connection from the processor to a bidirectional registered transceiver in an external agent. For those applications, the SysAD bus has only two possible drivers, the processor and the external agent. However, certain applications may wish to add additional drivers and receivers to the SysAD bus, and allow transmissions to take place over the SysAD bus that the processor is not involved in. To accomplish this the external agent must coordinate the usage of the SysAD bus using the arbitration handshake signals and the external null requests.

To implement an independent transmission on the SysAD bus that does not involve the processor, the external agent will request the SysAD bus to issue an external request. After the processor releases the system interface to slave state, the external agent may issue a scache release external null request to return ownership of the secondary cache to the processor, if the processor is being used with a secondary cache. The external agent may then allow the independent transmission to take place on the SysAD bus making sure that $\overline{\text{ValidIn}}$ is not asserted while the transmission is occurring. When the transmission is complete, the external agent will issue a system interface release external null request to return the system interface to master state.

4.8 Cycle Counts for System Interface Interactions

The V4400MC processor specifies minimum and maximum cycle counts for various processor transactions and for the processor's response time to external requests to facilitate system design with the V4400MC. Processor requests themselves are constrained by the system interface request protocol and the cycle counts for such requests can be determined by examining the protocol. The spacing between requests within a cluster, the waiting period for the processor to release the system interface to slave state in response to an external request, and the response time for an external request that requires a response is variable and subject to minimum and maximum cycle counts. The remainder of this section will describe and tabulate the minimum and maximum cycle counts for these system interface interactions.

The minimum and maximum number of unused cycles between the requests within a cluster is a function of processor internal activity. The minimum number of unused cycles separating requests within a cluster is zero, the requests may be adjacent. The maximum number of unused cycles separating requests within a cluster varies depending on the requests that form the cluster. The minimum and maximum number of unused cycles separating requests within a cluster is summarized in Table 4-5 **Unused Cycles Separating Requests Within a Cluster**.

Table 4-5 Unused Cycles Separating Requests Within a Cluster

From Processor Request	To Processor Request	Minimum Unused Cycles	Maximum Unused Cycles
Read	Update	0	2
Read	Write	0	2
Update	Write	0	2

The number of cycles the processor may wait to release the system interface to slave state for an external request will be referred to as the **release latency**. The release latency is a function of processor internal activity and processor request activity. The processor will release the system interface to accept an external request under the conditions described above. When no processor requests are in progress internal activity, such as refilling the primary cache from the secondary cache, may cause the processor to wait some number of cycles before releasing the system interface. Release latency will be considered in three categories:

- (1) release latency when the external request signal is asserted during the cycle two cycles before the last cycle of a processor request or two cycles before the last cycle of the last request in a cluster.
- (2) release latency when the external request signal is not asserted during a processor request or cluster, or asserts during the last cycle of a processor request or cluster.
- (3) release latency when the processor does an uncompelled change to slave state.
 - (a) Read with Write forthcoming (with respect to the last write data cycle).
 - (b) Read.

The minimum and maximum release latency for requests that fall into categories (1), (2) and (3) above is summarized in Table 4-6 **Release Latency for Category (1), (2) and (3) External Requests**.

Table 4-6 Release Latency for Category (1), (2) and (3) External Requests

<u>Category</u>	<u>Minimum PCycles</u>	<u>Maximum PCycles</u>
(1)	4	6
(2)	4	24
(3a)	0	Note
(3b)	0	Note

Note The Maximum Release Latency for;

(3a) Read with write forthcoming = $t_{D,s}$

- + four- or eight-word SCache Write Cycle Time
(depending on PCache size)
- + four-word SCache Write Cycle Time
- + SCline size (word)
- + 16

(3b) Read = four-word SCache Write Cycle Time

+ 4

The number of cycles the processor may take to respond to an external request that requires a response, that is, an external read request, intervention request, or snoop request will be referred to as the intervention response latency, external read response latency, or snoop response latency respectively. The number of cycles of latency is the number of unused cycles between the address cycle of the request and the first data cycle of the response. Intervention response latency and snoop response latency is a function of processor internal activity and secondary cache access time. The minimum and maximum intervention response latency and snoop response latency as a function of secondary cache access time is summarized in Table 4-7 **Intervention Response Latency and Snoop Response Latency**. External read response latency is purely a function of processor internal activity. The minimum and maximum external read response latency is summarized in Table 4-8 **External Read Response Latency**.

Table 4-7 Intervention Response Latency and Snoop Response Latency

<u>Maximum SCache Access</u>	<u>Intervention Response Latency</u>		<u>Snoop Response Latency</u>	
	<u>Minimum PCycles</u>	<u>Maximum PCycles</u>	<u>Minimum PCycles</u>	<u>Maximum PCycles</u>
<= 4 PCycles	6	26	6	26
5 – 6 PCycles	8	28	8	28
7 – 8 PCycles	10	30	10	30
9 – 10 PCycles	12	32	12	32
11 – 12 PCycles	14	34	14	34

Table 4-8 External Read Response Latency

	<u>Minimum PCycles</u>	<u>Maximum PCycles</u>
External Read Response Latency	4	4

4.9 System Interface Syntax

System interface commands specify the precise nature and attributes of any system interface request during the address cycle for the request. System interface data identifiers specify the attributes of a data element transmitted during a system interface data cycle. The following sections describe the syntax, that is the bitwise encoding, of system interface commands and data identifiers.

4.9.1 System Interface Command and Data Identifier Syntax

System interface commands and data identifiers are encoded in nine bits and transmitted from the processor to an external agent or from an external agent to the processor on the SysCmd bus during address and data cycles. Bit eight of the SysCmd bus determines whether the current contents of the SysCmd bus is a command or a data identifier and therefore whether the current cycle is an address cycle or a data cycle. For system interface commands SysCmd(8) must be de-asserted (0). For system interface data identifiers SysCmd(8) must be asserted (1).

For system interface commands and data identifiers associated with external requests, reserved bits and reserved fields in the command or data identifier should be de-asserted, that is set to one (1) or all ones respectively. For system interface commands and data identifiers associated with processor requests, reserved bits and reserved fields in the command or data identifier are undefined.

(1) System Interface Command Syntax

This section will define the encoding of the SysCmd bus for system interface commands. A common encoding is used for all system interface commands. SysCmd(8) must be de-asserted (0) for all system interface commands.

For all system interface commands SysCmd(7:5) specify the system interface request type which may be read, write, null, invalidate, update, intervention, or snoop. The encoding of SysCmd(7:5) for system interface commands is illustrated in Table 4-9 Encoding of SysCmd(7:5) for System Interface Commands.

Table 4-9 Encoding of SysCmd(7:5) for System Interface Commands

<u>SysCmd(7:5)</u>	<u>Command</u>
0	Read Request.
1	Read Request, Write Request forthcoming.
2	Write Request.
3	Null Request.
4	Invalidate Request.
5	Update Request.
6	Intervention Request.
7	Snoop Request.

For read requests, the remainder of the SysCmd bus specifies the attributes of the read. SysCmd(4:3) encode block, coherency, and exclusivity attributes for the read. A read request with a write request forthcoming cannot be a double word, word, or partial word read. For both coherent and noncoherent block reads SysCmd(2) specifies whether the address of the cache line being replaced by this read request is being retained in the link address register and SysCmd(1:0) encode the block size for the read. For double word, word, or partial word reads SysCmd(2:0) encode the size of the read data in bytes. The encoding of SysCmd(4:3) for read commands is shown in Table 4-10 Encoding of SysCmd(4:3) for Read Requests. The encoding of SysCmd(2:0) for block reads, or double word, word, or partial word reads is shown in Table 4-11 Encoding of SysCmd(2:0) for Block Read Requests, and Table 4-12 Encoding of SysCmd(2:0) for Double Word, Word, or Partial Word Read Requests respectively.

Table 4-10 Encoding of SysCmd(4:3) for Read Requests

<u>SysCmd(4:3)</u>	<u>Read attributes.</u>
0	Coherent block read.
1	Coherent block read, exclusivity requested.
2	Noncoherent block read.
3	Double word, single word, or partial word read.

Table 4-11 Encoding of SysCmd(2:0) for Block Read Requests

<u>SysCmd(2)</u>	<u>Link address retained indication.</u>
0	Link address not retained.
1	Link address retained.
<u>SysCmd(1:0)</u>	<u>Read block size.</u>
0	Four words.
1	Eight words.
2	Sixteen words.
3	Thirty-two words.

Table 4-12 Encoding of SysCmd(2:0) for Double Word, Word, or Partial Word Read Requests

<u>SysCmd(2:0)</u>	<u>Read data size.</u>
0	One byte valid. (Byte).
1	Two bytes valid. (Half Word).
2	Three bytes valid.
3	Four bytes valid. (Word).
4	Five bytes valid.
5	Six bytes valid.
6	Seven bytes valid.
7	Eight bytes valid. (Double Word).

For write requests, the remainder of the SysCmd bus specifies the attributes of the write. SysCmd(4:3) encode block attributes for the write. For block writes SysCmd(2) specifies whether the cache line associated with the write request will be replaced or retained after the write is completed and SysCmd(1:0) encode the block size for the write. For double word, word, or partial word writes SysCmd(2:0) encode the size of the write data in bytes. The encoding of SysCmd(4:3) for write commands is shown in Table 4-13 Encoding of SysCmd(4:3) for Write Requests. The encoding of SysCmd(2:0) for block writes or double word, word, or partial word writes is shown in tables Table 4-14 Encoding of SysCmd(2:0) for Block Write Requests and Table 4-15 Encoding of SysCmd(2:0) for Double Word, Word, or Partial Word Write Requests respectively.

Table 4-13 Encoding of SysCmd(4:3) for Write Requests

<u>SysCmd(4:3)</u>	<u>Write attributes.</u>
0	Reserved.
1	Reserved.
2	Block write.
3	Double word, single word, or partial word write.

Table 4-14 Encoding of SysCmd(2:0) for Block Write Requests

<u>SysCmd(2)</u>	<u>Cache line replacement attributes.</u>
0	Cache line replaced.
1 ^{Note}	Cache line retained.
<u>SysCmd(1:0)</u>	<u>Write block size.</u>
0	Four words.
1	Eight words.
2	Sixteen words.
3	Thirty-two words.
Note Valid only when Hit_Write_Back operation is set	

Table 4-15 Encoding of SysCmd(2:0) for Double Word, Word, or Partial Word Write Requests

<u>SysCmd(2:0)</u>	<u>Write data size.</u>
0	One byte valid. (Byte).
1	Two bytes valid. (Half Word).
2	Three bytes valid.
3	Four bytes valid. (Word).
4	Five bytes valid.
5	Six bytes valid.
6	Seven bytes valid.
7	Eight bytes valid. (Double Word).

Processor null write requests, system interface release external null requests, and scache release external null requests all use the null request command. For processor null requests, SysCmd(4:3) specifies that this is a null write request. For external null requests, SysCmd(4:3) specifies whether this is a system interface release null request or a scache release null request. The encoding of SysCmd(4:3) for processor null requests is shown in Table 4-16 Encoding of SysCmd(4:3) for Processor Null Requests. The encoding of SysCmd(4:3) for external null requests is shown in Table 4-17 Encoding of SysCmd(4:3) for External Null Requests.

Table 4-16 Encoding of SysCmd(4:3) for Processor Null Requests

<u>SysCmd(4:3)</u>	<u>Null attributes.</u>
0	Null write.
1	Reserved.
2	Reserved.
3	Reserved.

Table 4-17 Encoding of SysCmd(4:3) for External Null Requests

<u>SysCmd(4:3)</u>	<u>Null attributes.</u>
0	System interface release.
1	Secondary cache release.
2	Reserved.
3	Reserved.

For invalidate and update requests SysCmd(4) is used by external requests to indicate that this external request is in conflict with an unacknowledged processor invalidate or update request, canceling the invalidate or update. SysCmd(4) is reserved for processor invalidate and update requests. SysCmd(3) is used by processor requests to specify whether the update is potential or compulsory. SysCmd(3) is used by external update requests to indicate whether this update request will change the state of the updated cache line to shared, or leave the state of the updated cache line unchanged. SysCmd(2:0) specifies the size of the data element in bytes for updates. The encoding of SysCmd(4:0) for processor update requests is shown in Table 4-18 Encoding of SysCmd(4:0) for Processor Update Requests. The encoding of SysCmd (4:0) is reserved for processor invalidate requests. The encoding of SysCmd(4:0) for external invalidate and update requests is shown in Table 4-19 Encoding of SysCmd(4:0) for External Invalidate or Update Requests.

Table 4-18 Encoding of SysCmd(4:0) for Processor Update Requests

<u>SysCmd(4)</u>	<u>Reserved.</u>
<u>SysCmd(3)</u>	<u>Update type.</u>
0	Compulsory
1	Potential
<u>SysCmd(2:0)</u>	<u>Update data size.</u>
0	One byte valid (Byte).
1	Two bytes valid (Half Word).
2	Three bytes valid.
3	Four bytes valid (Word).
4	Five bytes valid.
5	Six bytes valid.
6	Seven bytes valid.
7	Eight bytes valid (Double Word).

Table 4-19 Encoding of SysCmd(4:0) for External Invalidate or Update Requests

<u>SysCmd(4)</u>	<u>Processor unacknowledged invalidate or update cancellation.</u>
0	Invalidate or update cancelled.
1	No cancellation.
<u>SysCmd(3)</u>	<u>Update cache state change attributes.</u>
0	Cache state changed to Shared.
1	No change to cache state.
<u>SysCmd(2:0)</u>	<u>Update data size.</u>
0	One byte valid (Byte).
1	Two bytes valid (Half Word).
2	Three bytes valid (Tri-Byte).
3	Four bytes valid (Word).
4	Five bytes valid (Quinti-Byte).
5	Six bytes valid (Sexti-Byte).
6	Seven bytes valid (Septi-Byte).
7	Eight bytes valid (Double Word).

4

For intervention and snoop requests SysCmd(4) is used to indicate that this external request is in conflict with an unacknowledged processor invalidate or update request, canceling the invalidate or update. The processor will never issue an intervention or snoop request. SysCmd(3) is the data response on dirty bit for intervention requests and is reserved for snoop requests. If the data response on dirty bit is de-asserted (0) the processor will return the contents of the cache line in response to an intervention request if the line is found in state dirty exclusive or dirty shared. If the data response on dirty bit is asserted (1) the processor will return the contents of the cache line in response to an intervention request if the line is found in state clean exclusive or dirty exclusive. For both snoop and intervention requests, SysCmd(2:0) specify a cache state change function to be applied to the cache line automatically with respect to the intervention or snoop response.

The encoding of SysCmd(4:0) for intervention requests is shown in Table 4-20 Encoding of SysCmd(4:0) for Intervention Requests. The encoding of SysCmd(4:0) for snoop requests is shown in Table 4-21 Encoding of SysCmd(4:0) for Snoop Requests.

Table 4-20 Encoding of SysCmd(4:0) for Intervention Requests

SysCmd(4)	<u>Processor unacknowledged invalidate or update cancellation.</u>
0	Invalidate or update cancelled.
1	No cancellation.
SysCmd(3)	<u>Data response on dirty bit.</u>
0	Return cache line data if in state dirty exclusive or dirty shared
1	Return cache line data if in state clean exclusive or dirty exclusive
SysCmd(2:0)	<u>Cache state change function.</u>
0	No change to cache state.
1	If cache state clean exclusive, change to shared, otherwise no change to cache state.
2	If cache state clean exclusive or shared, change to invalid, otherwise no change to cache state.
3	If cache state clean exclusive, change to shared or if cache state dirty exclusive, change to dirty shared, otherwise no change to cache state.
4	If cache state clean exclusive, dirty exclusive, or dirty shared, change to shared, otherwise no change to cache state.
5	Change to invalid regardless of current cache state.
6	Reserved.
7	Reserved.

Table 4-21 Encoding of SysCmd(4:0) for Snoop Requests

SysCmd(4)	<u>Processor unacknowledged invalidate or update cancellation.</u>
0	Invalidate or update cancelled.
1	No cancellation.
SysCmd(3)	<u>Reserved.</u>
SysCmd(2:0)	<u>Cache state change function.</u>
0	No change to cache state.
1	If cache state clean exclusive, change to shared, otherwise no change to cache state.
2	If cache state clean exclusive or shared, change to invalid, otherwise no change to cache state.
3	If cache state clean exclusive, change to shared or if cache state dirty exclusive, change to dirty shared, otherwise no change to cache state.
4	If cache state clean exclusive, dirty exclusive, or dirty shared, change to shared, otherwise no change to cache state.
5	Change to invalid regardless of current cache state.
6	Reserved.
7	Reserved.

(2) System Interface Data Identifier Syntax

This section will define the encoding of the SysCmd bus for system interface data identifiers. A common encoding is used for all system interface data identifiers. SysCmd(8) must be asserted (1) for all system interface data identifiers. System interface data identifiers have two formats, one for coherent data and a second for noncoherent data format. Data associated with processor block write requests and processor double word, word, or partial word write requests is noncoherent. Data associated with processor update requests is also noncoherent. Data returned in response to a processor coherent block read request is coherent while data returned in response to a processor noncoherent block read request or a processor double word, word, or partial word read request is noncoherent. Data associated with external update requests is noncoherent. Data associated with external write requests is noncoherent. Data returned in response to an external intervention request is coherent.

For both coherent and noncoherent data identifiers, both processor and external, SysCmd(7) marks the data element as the last data element, and SysCmd(6) indicates whether the data is response data or not. Response data is data returned in response to a read request or an intervention request. SysCmd(5) is the good data bit and indicates whether the data element is error free or not. Erroneous data contains an uncorrectable error. Erroneous data returned to the processor will cause a processor bus error. The processor will deliver data with the good data bit de-asserted when a primary parity error is detected for a transmitted data item. A secondary cache data ECC error can be detected by comparing the values transmitted on SysAD and SysADC. For external data identifiers, both coherent and noncoherent, SysCmd(4) indicates to the processor whether to check the data and check bits for this data element and SysCmd(3) is reserved. For processor data identifiers, both coherent and noncoherent, SysCmd(4:3) are reserved.

For coherent data identifiers SysCmd(2:0) indicate a cache state for the data. The cache state will provide the cache state with which to load the cache line for responses to processor coherent read requests. The cache state will indicate the cache state in which the line was found for data associated with the response to an external intervention request or for the data cycle issued in response to an external snoop request. For noncoherent data identifiers SysCmd(2:0) is reserved.

The encoding of SysCmd(7:3) for processor data identifiers is illustrated in Table 4-22 Encoding of SysCmd(7:3) for Processor Data Identifiers. The encoding of SysCmd(7:3) for external data identifiers is illustrated in Table 4-23 Encoding of SysCmd(7:3) for External Data Identifiers. The encoding of SysCmd(2:0) for coherent data identifiers is illustrated in Table 4-24 Encoding of SysCmd(2:0) for Coherent Data Identifiers.

Table 4-22 Encoding of SysCmd(7:3) for Processor Data Identifiers

<u>SysCmd(7)</u>	<u>Last data element indication.</u>
0	Last data element.
1	Not the last data element.
<u>SysCmd(6)</u>	<u>Response data indication.</u>
0	Data is response data.
1	Data is not response data.
<u>SysCmd(5)</u>	<u>Good data indication.</u>
0	Data is error free.
1	Data is erroneous.
<u>SysCmd(4:3)</u>	<u>Reserved.</u>

Table 4-23 Encoding of SysCmd(7:3) for External Data Identifiers

<u>SysCmd(7)</u>	<u>Last data element indication.</u>
0	Last data element.
1	Not the last data element.
<u>SysCmd(6)</u>	<u>Response data indication.</u>
0	Data is response data.
1	Data is not response data.
<u>SysCmd(5)</u>	<u>Good data indication.</u>
0	Data is error free.
1	Data is erroneous.
<u>SysCmd(4)</u>	<u>Data checking enable.</u>
0	Check the data and check bits.
1	Don't check the data and check bits.
<u>SysCmd(3)</u>	<u>Reserved.</u>

Table 4-24 Encoding of SysCmd(2:0) for Coherent Data Identifiers

<u>SysCmd(2:0)</u>	<u>Cache state.</u>
0	Invalid.
1	Reserved.
2	Reserved.
3	Reserved.
4	Clean Exclusive.
5	Dirty Exclusive.
6	Shared.
7	Dirty Shared.

4.10 System Interface Addresses

System interface addresses are full 36 bit physical addresses presented on the least significant 36 bits (bits 35 through 0) of the SysAD bus during address cycles. The remaining bits of the SysAD bus are unused during address cycles. Addresses associated with double word, word, or partial word transactions, i.e. double word, word, or partial word read and write requests and update requests, are aligned for the size of the data element. Specifically, for double word requests, the low order three bits of the address will be zero, for word requests, the low order two bits of the address will be zero, and for half-word requests, the low order bit of the address will be zero. For byte, tri-byte, quinti-byte, sexti-byte and septi-byte requests the address provided will be a byte address. For further details on addresses supplied by the V_R4400MC processor on double word, word, and partial word accesses see V_R44000, V_R4400 USER'S MANUAL—ARCHITECTURE.

Addresses associated with block requests are aligned to double word boundaries; that is the low order three bits of the address will be zero. The order in which data is returned in response to a processor block read request can be programmed via the boot time mode control interface to sequential ordering or sub-block ordering. If sequential ordering is enabled the processor will always deliver the address of the double word at the beginning of the block on a block read request. An external agent must return the block of data sequentially starting at the beginning of the block. If sub-block ordering is enabled the processor will deliver the address of the double word within the block that it wants returned first. An external agent must return the block of data using sub-block ordering starting with the addressed double word. For further details on sub-block ordering see **APPENDIX A SUB-BLOCK ORDERING**. Only a V_R4400MC with a secondary cache may be programmed to use sequential ordering.

For block write requests, the V_R4400MC processor will always deliver the double word address of the double word at the beginning of the block and deliver data beginning with the double word at the beginning of the block and progressing sequentially through the double words that form the block.

4

4.11 Processor Internal Address Map

External reads and writes to the V_R4400MC processor are provided to access processor internal resources that may be of interest to an external agent. However, the V_R4400MC does not contain any resources that are readable with an external read request. The V_R4400MC will return a bus error response to any external read request. The only writable resource in the V_R4400MC is the processor interrupts.

The processor decodes bits 6:4 of the address associated with an external read or write request to determine which processor internal resource is the target of the request. The only processor internal resource available for access by an external request is the interrupt resource, and it is only accessible via an external write request. The interrupt resource is accessed via an external write request with an address of 000 on bits 6:4 of the SysAD bus. See **CHAPTER 9 PROCESSOR INTERRUPTS** for further details on external writes to the interrupt resource.

4.12 Coherence Conflicts

The V_R4400MC processor will both issue processor coherence requests and accept external coherence requests. Processor coherence requests are processor coherent read requests, invalidate requests, and update requests. External coherence requests are external invalidate, update, snoop and intervention requests. Because of the overlapped nature of the system interface it is possible for processor coherence requests and external coherence requests to conflict. That is, it is possible for an external coherence request to reference an address that targets the same cache line as a pending processor read request or an unacknowledged processor invalidate or update request. The processor does not contain comparators to detect such conflicts. The processor uses the secondary cache as the single point of reference to determine the coherency actions it will take and only checks the state of the secondary cache at specific times.

For pending processor coherent read requests conflicting external requests cannot effect the behavior of the processor. The processor will only issue a read request for a particular cache line if it does not have a copy of that cache line. Therefore, any external coherence request that targets a cache line that is also the target of a pending processor coherent read request will not find the line present in the cache. External coherence requests do not change the state of the cache unless the cache line they target is present. Since no change can be made to the state of the cache for the line that is the target of the pending processor read request, no external coherence requests can effect the read request. Therefore, external coherence requests that conflict with a pending processor coherent read request may be issued to the processor and will effectively be discarded by the processor.

For processor invalidate and update requests the cancellation mechanism is provided to signal conflicts to the processor. If a conflicting external coherence request is issued while a processor invalidate or update request has been issued but not yet acknowledged, an external agent may cancel the processor invalidate or update (provided that the update request is compulsory). This is accomplished by setting the cancellation bit in the command for the external coherence request. The processor upon receiving an external coherence request with the cancellation bit set will consider its invalidate or update request to be acknowledged and canceled, re-access the secondary cache and re-evaluate the cache state to determine the correct action. This may result in the invalidate or update request being re-issued, or it may result in the issue of a read request instead.

An external agent is only allowed to assert the cancellation bit with an external coherence request when a processor invalidate or compulsory update request is currently unacknowledged. If an external coherence request is issued with the cancellation bit set when there is no unacknowledged processor invalidate or update request the behavior of the processor is undefined. Processor potential update requests may not be canceled. Potential updates are always issued under processor read requests and become compulsory only after the response to the processor read request is returned to the processor in one of the shared states. If an external coherence request is issued with the cancellation bit set when a processor update request is still potential, in other words while a processor read request is currently pending, the behavior of the processor is undefined.

An external agent may issue an external coherence request that is in conflict with an unacknowledged processor invalidate or update request without setting the cancellation bit. In this case the V_A4400MC will be unaware of the conflict and will not re-evaluate the cache state to determine the correct action. It will simply wait for an acknowledge to its invalidate or update request as for any invalidate or update request. A system employing the V_A4400MC may not behave correctly under these circumstances.

The interactions between processor coherence requests and conflicting external coherence requests, tabulated by processor state, is summarized in **Table 4-25 Coherence Conflicts Summary (a)** and **Table 4-26 Coherence Conflicts Summary (b)**. The processor can be in one of the following states:

- (1) Idle, no processor transactions currently pending;
- (2) Read Pending, a processor coherent read request has been issued but the read response has not yet been received;
- (3) Potential Update Unacknowledged, a processor update request has been issued while a processor coherent read request is pending but has not yet been acknowledged and by definition the response to the coherent read request has not yet been received;
- (4) Invalidate or Update Unacknowledged, a processor invalidate or update request has been issued but has not yet been acknowledged and by definition there is not a processor coherent read request pending.

Table 4-25 Coherence Conflicts Summary (a)

<u>Processor State</u>	<u>Conflicting External Coherent Request</u>			
	<u>Invalidate</u>	<u>Invalidate with Cancel</u>	<u>Update</u>	<u>Update with Cancel</u>
Idle	NA	Undefined	NA	Undefined
Read Pending	OK	Undefined	OK	Undefined
Potential Update Unacknowledged	OK	Undefined	OK	Undefined
Invalidate or Update Unacknowledged	OK <i>Note</i>	OK	OK <i>Note</i>	OK
Note This may cause incorrect system operation and should not normally be allowed to occur.				

4

Table 4-26 Coherence Conflicts Summary (b)

<u>Processor State</u>	<u>Conflicting External Coherent Request</u>			
	<u>Intervention</u>	<u>Intervention with Cancel</u>	<u>Snoop</u>	<u>Snoop with Cancel</u>
Idle	NA	Undefined	NA	Undefined
Read Pending	OK	Undefined	OK	Undefined
Potential Update Unacknowledged	OK	Undefined	OK	Undefined
Invalidate or Update Unacknowledged	OK <i>Note</i>	OK	OK <i>Note</i>	OK
Note This may cause incorrect system operation and should not normally be allowed to occur.				

4.13 System Implications of Coherence Conflicts

The constraints that the V_A4400MC processor places on the handling of conflicting coherency transactions has certain implications for the design of a system employing the V_A4400MC. This section will consider, as an example, a particular snoop, split-read, bus based system and the requirements for that system to correctly handle coherence conflicts.

4.13.1 System Model

The system model consists of the following components:

- (1) Four processor subsystems, each consisting of a V_A4400MC processor, a secondary cache, and an external agent. The agent communicates with the V_A4400MC, accepting processor requests and issuing external requests, and with the system bus likewise issuing and receiving bus requests.

(2) A memory subsystem that communicates with main memory and the system bus.

(3) A system bus with the following characteristics:

- It is a multiple master, request based, arbitrated bus in which an agent that wishes to perform a transaction on the bus must request the bus and wait for global arbitration logic to supply a grant signal before assuming mastership of the bus. Once mastership has been granted, the agent may begin a transaction.
- It supports a read transaction, read exclusive transaction, write transaction, and invalidate transaction.
- It is a split-read bus in that independent transactions may occur on the bus between a read request from a particular agent and the return of data by the target of the read request. The return of data by the target of the read request will be referred to as the read response.
- It is a snoopy bus in that all agents connected to the bus must monitor all of the traffic on the bus to correctly maintain cache coherency.
- I/O is not considered in this system model.

4.13.2 Coherency Model

The goal, for purposes of this example, is to implement a simple write invalidate cache coherency protocol for this system model that will maintain consistency between all of the caches in the system and main memory. Attention will be focused on the interactions between the system bus and the V_R4400MC and the handling of coherence conflicts that arise in this system.

The coherency model for the system is as follows: All pages in the system are maintained either with the noncoherent coherency attribute or with the sharable coherency attribute. The handling of noncoherent data will not be considered. Using the coherency attribute sharable allows data to be shared between the four caches in the system with a write invalidate cache coherency protocol. The secondary cache states used are invalid, shared, clean exclusive, and dirty exclusive. The secondary cache state dirty shared is not used in this coherency model.

When a processor misses in both caches on a load it issues a read request. The external agent will translate this to a read request on the bus. The returned data may be loaded in either the state clean exclusive or shared based on a shared indication returned on the bus with the read response. The shared indication is supplied by the external agents that are a part of the other three processor subsystems based on the result of an intervention request to the processor for the cache line of interest. When a processor misses in both caches on a store it issues a read request desiring exclusivity which is translated to a read exclusive on the bus and the data is loaded in the state dirty exclusive. When a processor hits in the cache on a store to shared data it issues an invalidate request which must be forwarded to the system bus before the store can be completed.

When an external agent observes a coherent read request on the system bus it does not take any immediate action, rather the external agent will issue an intervention request to the processor for the read request during the read response associated with the read request. This is referred to as a **response complete read model**; that is, the read is treated as complete only after the read response has occurred. This model requires that cache interrogation for a read must not occur until the read response occurs, as described, in order to maintain consistency. At the end of the read response each external agent will supply an indication on the bus of whether it was able to obtain the state of the cache line that is the target of the read via an intervention request; if it were able to obtain the state of the cache line, then it will indicate either shared or **takeover**. Takeover occurs when an external agent discovers that its processor has a copy of the cache line that is the target of the read in the state dirty exclusive. If any external agent is unable to obtain the state of the cache line that is the target of the read because it is unable to initiate an intervention request, the read response will be extended until all external agents have obtained the state of the cache line from their processors.

The response from an external agent at the end of a read response depends on whether the read request was an ordinary read request or a read exclusive request. For an ordinary read request an external agent will indicate shared at the end the read response if it finds that its processor has a copy of the requested cache line in the state clean exclusive or shared. If the current state of the cache line is clean exclusive the external agent will cause the processor to change the state of the cache line to shared. An external agent will indicate both shared and takeover at the end of a read response to an ordinary read request if it finds that its processor has a copy of the requested cache line in the state dirty exclusive. Having indicated takeover, the external agent will supply the contents of the cache line, returned by the processor in response to its intervention request, over the bus to the read requester, and cause the processor to change the state of the cache line to shared. At the same time the cache line is supplied to the read requester, it will also be written back to memory.

For a read exclusive request an external agent will never indicate shared at the end of the read response, regardless of the state its processor has the cache line in. If the current state of the cache line is clean exclusive or shared, the external agent will cause the state of the cache line to be changed to invalid. If the current state of the cache line is dirty exclusive, the external agent will indicate takeover but not shared. Having indicated takeover, the external agent will supply the contents of the cache line over the bus to the read requester, and cause the processor to change the state of the cache line to invalid. At the same time the cache line is supplied to the read requester, it will also be written back to memory.

An invalidate request will be considered complete as soon as it appears on the system bus. When an external agent observes an invalidate request on the system bus it must react as if the invalidate has changed the state of all caches at that instant.

An external agent will take no action in response to the appearance of a write request on the bus.

4.13.3 Handling Coherence Conflicts

(1) Coherent Read Conflicts

Coherence conflicts can be examined based on the current state of the processor. In particular, the processor may have a coherent read request pending, or it may have an invalidate request unacknowledged, or it may not have any requests pending or unacknowledged. Note that the read exclusive transaction on the system bus guarantees that the requested cache line will return in an exclusive state.

External coherence requests that conflict with pending processor coherent read requests may be issued to the processor without effecting the processor's behavior. Therefore, in this simple system model no conflict detection will be performed by the external agent for processor coherent read requests. If an external intervention request or invalidate request is forwarded to the processor that is in conflict with a pending processor coherent read request it will not effect the processor's cache since the target cache line is guaranteed to be absent from the cache. The processor will effectively discard a conflicting external intervention request, responding with an invalid indication for the target cache line. Similarly, the processor will discard a conflicting external invalidate request since the target cache line is not present and therefore already invalid.

In a system model similar to the one described conflict detection could be provided for pending processor coherent read requests. In this case, when the external agent sees a read response on the bus that conflicts with a pending processor coherent read request, it will not issue an intervention request to the processor. Rather, it will simply react as if an intervention request has been completed and the cache line is not present in the processor's cache. Similarly, when the external agent sees an invalidate request on the bus that conflicts with a pending processor coherent read request, it will not forward the invalidate request to the processor since the target cache line is known to be absent from the processor's cache. Using this scheme for conflict detection on processor coherent read requests might slightly reduce the number of external intervention and invalidate requests issued to the processor. However, since the intervention and invalidate requests that would otherwise be issued to the processor would not result in any state modification within the processor, conflict detection for processor coherent read requests is not necessary.

(2) Invalidate Conflicts

From the time the processor has issued an invalidate request until the request has been acknowledged any external coherence request issued to the processor that is in conflict with the unacknowledged invalidate must include a cancellation. In this system, an acknowledge for the invalidate will be generated to the processor as soon as the invalidate is forwarded to the system bus. Therefore, while the external agent is waiting to acquire mastership of the system bus to forward an invalidate request the external agent must detect, via external comparators, any external coherence request that conflicts with the unacknowledged invalidate. If a conflict is detected the external agent must not forward the invalidate request to the system bus, rather, it must throw the invalidate request away and issue the conflicting external request to the processor with a cancellation.

If the response to a coherent read request conflicts with a waiting unacknowledged processor invalidate request appears on the bus the external agent will detect the conflict and will not forward the processor invalidate request to the bus. Rather, it will throw the processor invalidate request away and issue an intervention request to the processor that includes a cancellation. The processor will then re-evaluate its cache state and re-issue the invalidate request or issue a coherent read request instead.

If an invalidate request appears on the bus while the external agent has a processor invalidate request waiting and the external agent detects a conflict, the external agent will not forward the processor invalidate request to the bus. Rather, it will throw the processor invalidate request away and issue an external invalidate request to the processor that includes a cancellation. The processor will then re-evaluate its cache state and re-issue the invalidate request or issue a coherent read request instead.

(3) Write Conflicts

As soon as a write request has been issued to the external agent the external agent becomes responsible for the cache line. No coherence conflicts are possible with a processor write request, however, the external agent must manage ownership of the cache line while it is waiting to acquire mastership of the system bus to forward the write request. The external agent is responsible for the cache line from the time the issue cycle of the write request is accomplished until the write request is forwarded to the system bus.

If the response to a coherent read request conflicts with a waiting processor write request or with a processor write request that has issued and is transmitting data appears on the bus the external agent will detect the conflict and will not issue an intervention request to the processor. Rather, it will react as if an intervention request has been completed and the line is in the dirty exclusive state. The external agent will indicate takeover and supply the read data to the read requester itself without disturbing the processor. After providing the read data to the read requester the external agent must throw the write request away if the read request was a read exclusive. In fact, the external agent may throw the write request away for either type of read since processor supplied read data is also written back to memory.

It is not possible for an invalidate request or a write request that conflicts with a waiting processor write request to appear on the system bus since for a processor write request to be issued the state of the cache line must be dirty exclusive in that processor's cache.

CHAPTER 5 ERROR CHECKING AND CORRECTING (ECC)

The Va4400MC processor provides sixteen check bits for the secondary cache data bus, seven check bits for the secondary cache tag bus and eight check bits for the system interface address and data bus. The sixteen check bits for the secondary cache data bus are organized as eight check bits for the upper sixty-four bits of the data bus and eight check bits for the lower sixty-four bits of the data bus. In addition, a single check bit is provided for the system interface command bus.

The eight check bits for the system interface address and data bus provide either even byte parity or are generated in accordance with a single error correcting double error detecting (SECDED) code that also detects any three or four bit error in a nibble. The eight check bits for each half of the secondary cache data bus are always generated in accordance with the SECDED code.

The processor checks data using parity or the SECDED code as it passes from the system interface to the secondary cache and as it is moved from the secondary cache to the primary cache or to the system interface. The processor passes the check bits for data accessed from the secondary cache directly to the system interface without change as it checks it. The processor does not check data received from the system interface for external updates/external writes. It is possible to force the processor to not check data from the system interface for read responses using a bit in the data identifier. The processor does generate correct check bits for double word, word, or partial word data transmitted to the system interface. The processor does not check addresses received from the system interface, but does generate correct check bits for addresses transmitted to the system interface. The processor does not contain a data corrector, rather, the processor will take Cache Error Expection when an error is detected based on the data check bits. If the system interface is set in the parity mode, the processor indicates the ECC error of the secondary cache by driving a wrong party to the SysCmd signal. Software, in conjunction with an off processor data corrector, is responsible for correcting the data when the SECDED code is employed.

The seven check bits for the secondary cache tag bus are generated in accordance with a single error correcting double error detecting (SECDED) code that also detects any three or four bit error in a nibble. The processor generates check bits for the tag when it is written into the secondary cache and checks the tag whenever the secondary cache is accessed. The processor contains a corrector for the secondary cache tag. The tag corrector is not in-line for processor accesses due to primary cache misses. When a tag error is detected on a processor access due to a primary cache miss the processor will trap. Software, using the Va4400MC cache management primitives, will cause the tag to be corrected. When executing the cache management primitives, the processor uses the corrected tag to generate write back addresses and cache state. For external accesses the tag corrector is in-line; that is the response to external accesses will be based on the corrected tag. The processor will still trap on tag errors detected during external accesses to allow software to repair the contents of the cache if possible.

The check bit for the system interface command bus provides even parity over the nine bits of the system interface command bus. This parity bit is generated correctly when the system interface is in master state, but is not checked when the system interface is in slave state. Even when returning data of less than 64 bits, such as a byte, in a read response, generate check bits for 64-bit data.

The buses that are covered by check bits and their contents and whether they are checked or not for various processor internal and external transactions is summarized in Table 5-1 Error Checking and Correcting Summary for Internal Transactions and Table 5-2 Error Checking and Correcting Summary for External Requests.

Table 5-1 Error Checking and Correcting Summary for Internal Transactions (1/2)

<u>Bus</u>	<u>Secondary Cache to Primary Cache</u>	<u>Primary Cache to Secondary Cache</u>	<u>Processor Read Request</u>	<u>Processor Write Request</u>
SCData	Checked, Cache Error Exception on Error	Primary Cache Parity checked, Cache Error Exception on Error	From System Interface	Not Checked
SCDChk	Checked, Cache Error Exception on Error	Generated	NA	NA
SCTag, SCTChk	Checked, Not Corrected, Cache Error Exception on Error	NA	NA	NA
Output of SysAD (address), SysCmd, SysADC (address), SysCmdP	NA	NA	Generated	Generated
Input of SysAD (address), SysCmd, SysADC (address), SysCmdP	NA	NA	Not Checked (Indicated by Fault pin on Error)	NA
SysAD (data)	NA	NA	Checked, Cache Error Exception on Error ^{Note}	From Processor
SysADC (data)	NA	NA	Checked, Cache Error Exception on Error ^{Note}	Generated
Note If the ERL bit of the status register is 1, an exception occurs.				

Table 5-1 Error Checking and Correcting Summary for Internal Transactions (2/2)

<u>Bus</u>	<u>Store to Shared Cache Line</u>	<u>Cache Instruction</u>	<u>Secondary Cache Miss</u>	<u>Secondary Cache Write to System Int</u>
SCData	NA	Checked on Write Back, Cache Error Exception on Error	From System Interface, unchanged	Checked, Cache Error Exception on Error
SCDChk	NA	Checked on Write Back, Cache Error Exception on Error	From System Interface, unchanged	Checked, Cache Error Exception on Error
SCTag, SCTChk	Checked on read part of RMW, Cache Error Exception on Correcting Secondary Cache Tag	Checked, Cache Error Exception on Correcting Secondary Cache Tag ^{Note 1}	Generated	Checked, not corrected, Cache Error Exception on Error
Output of SysAD (address), SysCmd, SysADC (address), SysCmdP	NA	Generated	Generated	Generated
Input of SysAD (address), SysCmd, SysADC (address), SysCmdP	From Processor	NA	Not Checked	NA
SysAD (data)	Generated	From Primary or Secondary Cache	Checked, Cache Error Exception on Error ^{Note 2}	From Secondary Cache
SysADC (data)	Generated	From Primary or Secondary Cache	Checked, Cache Error Exception on Error ^{Note 2}	From Secondary Cache

Notes 1. If the current cache instruction requires change of tag and write back.
2. If the ERL bit of the status register is 1, an exception occurs.

Table 5-2 Error Checking and Correcting Summary for External Requests (1/2)

<u>Bus</u>	<u>Read Request</u>	<u>Write Request</u>	<u>Invalidate Request</u>	<u>Update Request</u>
SCData	NA	NA	Not Checked	Checked on read part of RMW, Cache Error Exception on Error
SCDChk	NA	NA	Not Checked	Checked on read part of RMW ^{Note} , Cache Error Exception on Error, Generated on write part of RMW if written
SCTag, SCTChk	NA	NA	Checked on read part of RMW, Cache Error Exception on Error, Generated on write part of RMW if written	Checked on read part of RMW, Cache Error Exception on Error, Generated on write part of RMW if written
Output of SysAD (address), SysCmd, SysADC (address), SysCmdP	Generated	NA	NA	NA
Input of SysAD (address), SysCmd, SysADC (address), SysCmdP	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)
SysAD (data)	From processor	Checked, Cache Error Exception on Error	Not checked	Not checked (Indicated by $\overline{\text{Fault}}$ pin on Error)
SysADC (data)	From processor	Checked, Cache Error Exception on Error	Not checked	Not checked (Indicated by $\overline{\text{Fault}}$ pin on Error)
Note Only the double word being updated is checked.				

Table 5-2 Error Checking and Correcting Summary for External Requests (2/2)

<u>Bus</u>	<u>Intervention Request Data Returned</u>	<u>Intervention Request State Returned</u>	<u>Snoop Request</u>
SCData	Checked, Cache Error Exception on Error	Not Checked	Not Checked
SCDChk	Checked, Cache Error Exception on Error	Not Checked	Not Checked
SCTag, SCTChk	Checked and corrected on read part of RMW, Trap on Error, Generated on write part of RMW if written.	Checked and corrected on read part of RMW, Trap on Error, Generated on write part of RMW if written.	Checked and corrected on read part of RMW, Trap on Error, Generated on write part of RMW if written.
Ouput of SysAD (address), SysCmd, SysADC (address), SysCmdP	Generated	Generated	Generated
Input of SysAD (address), SysCmd, SysADC (address), SysCmdP	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)	Not Checked (Indicated by $\overline{\text{Fault}}$ pin on Error)
SysAD (data)	From Secondary Cache	NA	NA
SysADC (data)	From Secondary Cache	NA	NA

5.1 Single Error Correcting Double Error Detecting Codes

The ECC codes chosen for the processor's secondary cache data and secondary cache tag are single error correcting double error detecting codes that also detect three or four bit errors within a nibble. The 64 bit data code can detect three and four bit errors within a nibble. The 25 bit tag code was created using the patterns observed in the 64 bit data code.

The data code has the following properties:

Nibble is defined by the 4-bit vertical group shown in Tables 5-3 and 5-4.

- (1) It is a single error correcting, double error and three or four bit error within a nibble detecting code.
- (2) It provides 64 data bits protected by 8 check bits yielding 8 bit syndromes.
- (3) It is minimal in that each parity tree used to generate the syndrome has only 27 inputs, the minimum possible number.
- (4) It provides byte XORs of the data bits as part of the XOr trees used to build the parity generators. This allows picking byte parity out of the XOr trees that generate or check the code.
- (5) Single bit errors are indicated by syndromes that contain exactly 3 ones or by syndromes that contain exactly 5 ones in which bits 0-3 or bits 4-7 of the syndrome are all one. This makes it possible to decode the syndrome to find which data bit is in error with 4 input NAND gates, provided a pre-decode AND of bits 0-3 and bits 4-7 of the syndrome is available. For the check bits a full 8 bit decode of the syndrome is required.
- (6) Double bit errors are indicated by syndromes that contain an even number of ones.
- (7) Three bit errors within a nibble are indicated by syndromes that contain 5 ones in which bits 0-3 of the syndrome and bits 4-7 of the syndrome are not all one.
- (8) Four bit errors within a nibble are indicated by syndromes that contain 4 ones. Because this is an even number of ones, four bit errors within a nibble look like double bit errors.

The tag code has the following properties:

- (1) It is a single error correcting, double error and three or four bit error within a nibble detecting code.
- (2) It provides 25 data bits protected by 7 check bits yielding 7 bit syndromes.
- (3) It provides byte XORs of the data bits as part of the XOr trees used to build the parity generators. This allows picking byte parity out of the XOr trees that generate or check the code.
- (4) Single bit errors are indicated by syndromes that contain exactly 3 ones. This makes it possible to decode the syndrome to find which data bit is in error with 3 input NAND gates. For the check bits a full 7 bit decode of the syndrome is required.
- (5) Double bit errors are indicated by syndromes that contain an even number of ones.
- (6) Three bit errors within a nibble are indicated by syndromes that contain 5 ones or 7 ones.
- (7) Four bit errors within a nibble are indicated by syndromes that contain 4 ones or 6 ones. Because these are even numbers of ones, four bit errors within a nibble look like double bit errors.

The parity check matrices for the data ECC code and the tag ECC code specifying the distribution of data and check bits across nibbles are shown in Fig. 5-1 Parity Check Matrix for the Data ECC Code and Fig. 5-2 Parity Check Matrix for the Tag ECC Code.

Fig. 5-1 Parity Check Matrix for the Data ECC Code

[illegible]

Remarks 1. Data bit numbers are described vertically. For example, a_6 indicates bit 63.

2. Bit (7:0) of the Check Bit row correspond to the following Vn4400MC check bit signals:
SysAD(7:0), SCDCChk(15:8), SCDCChk(7:0)
3. Bit (63:0) of the Data Bit row correspond to the following Vn4400MC data bus signals:
SysAD(63:0), SysAD(127:64), SCData(63:0), SCData(127:64)

Fig. 5-2 Parity Check Matrix for the Tag ECC Code

Check Bit		0	12	34	56				
Data Bit		222 432	22 10	11 98	11 76	1111 5432	11 1098	7654	3210
Number of ones per row	11	. 1 . .	1 . . .	1 . . .	. . . 1	1111	1 . . .	1 . . .	1 . . .
	13	1 . . .	. 1 . .	. 1 . .	. . 1 .	1111	1111		. 1 . .
	10	. . 1 .	1 . . .	. . . 1	1 . . .		1111	. 1 . .	. . 1 .
	10	. 1 . .	. 1 . .	. . 1 .	. 1 . .	1 . . .	. 1 . .	1111	
	13	1 . . .	. . . 1	1 . . .	1 . . .	. 1 . .		1111	1111
	11	. . 1 .	. . 1 .	. 1 . .	. 1 . .	. . 1 .	. . 1 .	. . 1 .	1111
	14	1111	11 . .	11 . .	11 . .	. . . 1	. . . 1	. . . 1	. . . 1
Number of ones per column		3331	3311	3311	3311	3333	3333	3333	3333

- Remarks
1. Data bit numbers are described vertically. For example, $\begin{smallmatrix} 2 \\ 4 \end{smallmatrix}$ indicates bit 24.
 2. Bit (6:0) of the Check Bit row correspond to SCTChk (6:0).
 3. Bit (24:0) of the Data Bit row correspond to SCTag (24:0).

CHAPTER 6 BOOT TIME MODE CONTROL INTERFACE

Fundamental operational modes for the processor are initialized via the boot time mode control interface. The boot time mode control interface is a serial interface operating at a very low frequency to allow the initialization information to be kept in a low cost EPROM.

Immediately after the $V_{DD}Ok$ signal is asserted, the processor will read a serial bit stream of 256 bits to initialize all fundamental operational modes. After initialization is complete, the processor will continue to drive the serial clock output but no further initialization bits will be read.

6.1 Boot Time Mode Control Interface Signal Summary

ModeIn: (i) Serial boot mode data in.


ModeClock: (o) Serial boot mode data clock out at the MasterClock frequency divided by 256.

6

6.2 Boot Time Mode Control Interface Operation

While the $V_{DD}Ok$ signal is de-asserted, the **ModeClock** output will be held asserted. After the $V_{DD}Ok$ signal is asserted, the first bit in the initialization bit stream must be present at the **ModeIn** input. The processor will synchronize the **ModeClock** output at the time $V_{DD}Ok$ is asserted, and the first rising edge of the **ModeClock** will occur 256 MasterClock (MCik) cycles after $V_{DD}Ok$ is asserted. The processor will sample exactly 256 initialization bits from the **ModeIn** input one master clock cycle before the rising edge of the **ModeClock**. After each rising edge of the **ModeClock**, the next bit of the initialization bit stream must be presented at the **ModeIn** input.

Fig. 6-1 BTMC Interface Timing



6.3 Boot Time Modes

The correspondence between bits of the initialization bit stream and processor mode settings is illustrated in Table 6-1 **Boot Time Modes**. Bit 0 of the bit stream is the bit presented to the processor when $V_{DD}Ok$ is de-asserted.

Note the followings:

- (1) Selecting a reserved value results in undefined processor behavior.
- (2) Zeros must be scanned in for all reserved bits.

Table 6-1 Boot Time Modes (1/4)

Serial Bit	Value	Mode Setting
0	0	BlkOrder: Secondary Cache Mode block read response ordering.
	1	Sequential ordering.
1	0	Sub-block ordering.
	1	EIBParMode: Specifies nature of system interface check bus.
2	0	SECEDED error checking and correcting mode.
	1	Byte parity.
3	0	EndBit: Specifies byte ordering.
	1	Little Endian ordering.
4	0	Big Endian ordering.
	1	DSHMdDis: Dirty shared mode, enables transition to dirty shared state on processor update successful.
5:6	0	Dirty shared mode enabled.
	1-3	Dirty shared mode disabled.
7	0	NoSCMode: Specifies presence of secondary cache.
	1	V _R 4000SC, V _R 4400SC and V _R 4400MC.
8	0	V _R 4000PC and V _R 4400PC.
	1	SysPort: System Interface port width, bit 6 is the most significant.
9:10	0	64 bits.
	1-3	Reserved.
11:14	0	SC64BitMd: Secondary cache interface port width.
	1	128 bits.
15:16	0	Reserved.
	1	EISpltMd: Specifies secondary cache organization.
17:18	0	Secondary cache unified.
	1	Secondary cache split instruction vs. data.
19:20	0	SCBlkSz: Secondary cache line size, bit 10 is the most significant.
	1	4 words.
21:22	2	8 words.
	3	16 words.
23:24	0	32 words.
	1	XmitDatPat: System interface data rate, bit 14 is the most significant.
25:26	0	D
	1	DDx
27:28	2	DDxx
	3	DxDx
29:30	4	DDxxx
	5	DDxxxx
31:32	6	DxxDxx
	7	DDxxxxxx
33:34	8	DxxxDxxx
	9-15	Reserved.

Table 6-1 Boot Time Modes (2/4)

Serial Bit	Value	Mode Setting															
15:17		SysCkRatio: PClock to SClock divisor, frequency relationship between SClock, RClock, and TClock and PClock, bit 17 is the most significant.															
	0	Divide by 2.															
	1	Divide by 3.															
	2	Divide by 4.															
	3	Divide by 6.															
	4	Divide by 8.															
	5-7	Reserved.															
18		SIMasterMd: System Interface Master. Specifies the operation of the processor in Master/Checker mode combined with SCMasterMd(bit 42).															
		<table> <tr> <th>SIMasterMd</th><th>SCMasterMd</th><th>Mode</th></tr> <tr> <td>0</td><td>0</td><td>Complete Master</td></tr> <tr> <td>0</td><td>1</td><td>System Interface Master</td></tr> <tr> <td>1</td><td>0</td><td>Secondary Cache Master</td></tr> <tr> <td>1</td><td>1</td><td>Complete Listener</td></tr> </table>	SIMasterMd	SCMasterMd	Mode	0	0	Complete Master	0	1	System Interface Master	1	0	Secondary Cache Master	1	1	Complete Listener
SIMasterMd	SCMasterMd	Mode															
0	0	Complete Master															
0	1	System Interface Master															
1	0	Secondary Cache Master															
1	1	Complete Listener															
19		TimIntDis: Timer Interrupt disables timer interrupts, and the interrupt used by the timer becomes a general-purpose interrupt.															
	0	Timer Interrupt enabled.															
	1	Timer Interrupt disabled.															
20		PotUpdDis: Potential update enable allows potential update to be issued. Otherwise only compulsory updates are issued.															
	0	Potential updates enabled.															
	1	Potential updates disabled.															
21:24		TWrSup: Secondary cache write deassertion delay, TWrSup in PCycles, bit 24 is the most significant.															
	0-2	Undefined.															
	3-15	Number of PClock cycles; MIN. 3, MAX. 15.															
25:26		TWr2Dly: Secondary cache write assertion delay 2, TWr2Dly in PCycles, bit 26 is the most significant.															
	0	Undefined.															
	1-3	Number of PClock cycles; MIN. 1, MAX. 3.															
27:28		TWr1Dly: Secondary cache write assertion delay 1, TWr1Dly in PCycles, bit 28 is the most significant.															
	0	Undefined.															
	1-3	Number of PClock cycles; MIN. 1, MAX. 3.															

Table 6-1 Boot Time Modes (3/4)

Serial Bit	Value	Mode Setting
29	0	TWrRc: Secondary cache write recovery time, TWrRc in PCycles, either 0 or 1 cycle.
	1	0 cycle.
30:32	0-1	TDis: Secondary cache disable time, TDis in PCycles, bit 32 is the most significant.
	2-7	Undefined.
33:36	0-2	Number of PClock cycles; MIN. 2, MAX. 7.
	3-15	TRd2Cyc: Secondary cache read cycle time 2, TRdCyc2 in PCycles, bit 36 is the most significant.
37:40	0-3	Undefined.
	4-15	TRd1Cyc: Secondary cache read cycle time 1, TRdCyc1 in PCycles, bit 40 is the most significant.
41	0	NoMPmode: Multi-processor mode.
	1	Sets valid lines in the secondary cache to "Invalid" after a secondary cache miss. (Set after a write back if necessary) Set to 0 when using the multi-processor system. Does not set valid lines in the secondary cache to "Invalid" after a secondary cache miss (to improve performance)
42		SCMasterMd: Secondary Cache Master. Specifies the operation of the processor in Master/Checker mode combined with SIMasterMd(bit 18). See bit 18 for setting of these bits.
43:45	0	Reserved.
46 Note	0	Pkg179: Package type.
	1	Vr4400SC, Vr4400MC and Vr4400PC.
47:49	0	CycDivisor: This mode determines the clock divisor for the reduced-power mode.
	1	When the RP bit in the Status Register is set to one, the pipeline clock is divided by one of the following values. Bit 49 is the most significant.
50:52	2	Divide by 2
	3	Divide by 4
	4	Divide by 8
	others	Divide by 16
	4-7	Reserved.
	50:52	Drv0_50, Drv0_75, Drv1_00: Drive-off time. Drive the outputs out in (M x MasterClock) period. Drv0_50 (bit 50) is the least significant bit and Drv1_00 (bit 52) is the most significant bit. Refer to chapter 12.
	1	0.50 * MasterClock Period
	2	0.75 * MasterClock Period
	4	1.00 * MasterClock Period
	others	Reserved.

Note Be sure to set this bit to 0. If 1 is set, the behavior of the Vr4400MC is undefined.

Table 6-1 Boot Time Modes (4/4)

Serial Bit	Value	Mode Setting
53:56	Notes 1,2	InitP: Initial values for the state bits that determine the pull-down di/dt and switching speed of the output buffers. Bit 53 is the most significant.
	0	Fastest pull-down rate.
	1-14	Intermediate pull-down rates.
	15	Slowest pull-down rate.
57:60	Notes 1,2	InitN: Initial values for the state bits that determine the pull-up di/dt and switching speed of the output buffers. Bit 57 is the most significant.
	0	Slowest pull-up rate.
	1-14	Intermediate pull-up rates.
	15	Fastest pull-up rate.
61	Note 2	EnbIDPLL: Enables the negative feedback loop that controls switching speed of the output buffers only during ColdReset.
	0	Disable di/dt control mechanism.
	1	Enable di/dt control mechanism.
62	Note 2	EnbIDPLL: Enables the negative feedback loop that determines the di/dt and switching speed of the output buffers during ColdReset and during normal operation.
	0	Disable di/dt control mechanism.
	1	Enable di/dt control mechanism.
63		DsbIDPLL: Enables PLLs that match MasterIn and produce RClock, TClock, SClock and the internal clocks.
	0	Enable PLLs.
	1	Disable PLLs.
64		SRTriState: Controls when output-only pins are 3-stated.
	0	Only when $\overline{\text{ColdReset}}$ is asserted.
	1	When $\overline{\text{Reset}}$ or $\overline{\text{ColdReset}}$ are asserted.
65:255	0	Reserved.

Notes 1. Set InitP to 0 and InitN to 15 for fastest pull-up and pull-down rate.

Set InitP to 15 and InitN to 0 for slowest pull-up and pull-down rate.

2. Set InitP and InitN to the slowest pull-up and pull-down rate when the di/dt control is enabled by the bit EnbIDPLL or the bit EnbIDPLL.

CHAPTER 7 RESET SEQUENCE FOR THE V_A4400MC PROCESSOR

The V_A4400MC processor supports a multi-level reset sequence using the $\overline{V_{DD}Ok}$, $\overline{ColdReset}$, and $\overline{Reset}$ inputs.

7.1 Power-on Reset

For a power-on reset, the following sequence must be applied to the V_A4400MC:

- (1) Stable V_{DD} within the specified range from the power supply must be applied to the V_A4400MC continuously for at least 100 milliseconds along with a stable continuous system clock at the desired operational frequency of the V_A4400MC.
- (2) After at least 100 milliseconds of stable V_{DD} and MasterClock the $\overline{V_{DD}Ok}$ input to the V_A4400MC may be asserted. While $\overline{V_{DD}Ok}$ is de-asserted, all pins will be 3-stated with the exception of ModeClock. The assertion of $\overline{V_{DD}Ok}$ will cause the V_A4400MC to begin driving a clock on the ModeClock output and to read the boot time mode control serial data stream. For further details on mode initialization see CHAPTER 6 BOOT TIME MODE CONTROL INTERFACE. After the mode bits have been read in, the V_A4400MC will allow its internal phase locked loops to lock, stabilizing the processor internal clock, PClock, the SyncOut-SyncIn clock path and the master clock output, MasterOut. $\overline{V_{DD}Ok}$ must become asserted at least 100 milliseconds before the de-assertion of $\overline{ColdReset}$. If JTAG is not used, fix JTCK to low level. To use JTAG, input the clock to JTCK during reset sequence.
- (3) Once the boot time mode control serial data stream has been read by the V_A4400MC, the $\overline{ColdReset}$ input may be de-asserted. $\overline{ColdReset}$ must remain asserted for at least 100 milliseconds after the assertion of $\overline{V_{DD}Ok}$. $\overline{ColdReset}$ must be de-asserted synchronously with MasterClock. The state of SClock, TClock, and RClock is undefined until $\overline{ColdReset}$ is de-asserted. The de-asserted edge of $\overline{ColdReset}$ is used to synchronize the edges of SClock, TClock, and RClock, potentially across multiple processors in a multi-processor system.
- (4) De-assert $\overline{ColdReset}$ in synchronization with rising edges of SClock and TClock and the next rising edge of MasterClock. This signal is used to synchronize multiple processors in a multi-processor system. SClock, TClock, and RClock becomes stable 64 MasterClock cycles after $\overline{ColdReset}$ has been de-asserted.
- (5) After $\overline{ColdReset}$ is de-asserted and SClock, TClock and RClock have stabilized, $\overline{Reset}$ may be de-asserted to allow the V_A4400MC to begin to run. $\overline{Reset}$ must be held asserted for at least 64 MasterClock cycles after the de-assertion of $\overline{ColdReset}$. $\overline{Reset}$ must be de-asserted synchronously with MasterClock.

7.2 Cold Reset

A cold reset requires the same sequence as described above except that power is presumed to have been stable for a long time before the assertion of the reset inputs and the de-assertion of $\overline{V_{DD}Ok}$. $\overline{V_{DD}Ok}$ must be de-asserted for a minimum of 64 MasterClock cycles before re-asserting to begin the reset sequence.

7.3 Warm Reset

To effect a warm reset, the $\overline{Reset}$ input may be asserted synchronously with MasterClock and held asserted for at least 64 MasterClock cycles before being de-asserted synchronously with MasterClock. The processor internal clocks, PClock and SClock, and the system interface clocks, TClock and RClock will not be affected by a warm reset, and the boot time mode control serial data stream will not be read by the V_A4400MC on a warm reset.

7.4 Notes on Reset

After a power-on reset, cold reset, or warm reset, all processor internal state machines will be reset, and the processor will begin execution at the reset vector. The processor internal state is preserved during a warm reset, although the precise state of the caches will depend on whether a cache miss sequence has been interrupted by resetting the processor state machines.

ColdReset must be asserted when VddOk asserts. The behavior of the Vr4400MC is undefined if VddOk asserts while ColdReset is de-asserted.

Timing diagrams for reset sequence are illustrated in Fig.7-1, Fig.7-2, and Fig.7-3.

Fig. 7-1 Power-on Reset or Cold Reset




Fig. 7-2 Cold Reset





Fig. 7-3 Warm Reset



CHAPTER 8 MASTER/CHECKER MODE

Two V_R4400MC's can be configured as a lock-stepped pair to improve data integrity. For such configurations, outputs (and I/O busses during output) are checked between the two chips connected in parallel at bus cycle boundaries. Any discrepancies between them are reported through the $\overline{\text{Fault}}$ pin.

Two boot-time mode bits determine these Master/Checker configurations. For single-chip operations, these bits must be set to 00. This is the "Complete Master" mode, where the V_R4400MC's outputs operate normally.

There are two possible lock-step configurations. The simple case pairs a "Complete Master" V_R4400MC (with mode bits set to 00) with a "Complete Listener" V_R4400MC, with mode bits set to 11. The latter has output drivers disabled. Since it receives all the same inputs as the former, both V_R4400MC's operate identically. On all output cycles, it samples the signals on output and I/O busses and compares with what it expects. Any discrepancies are reported through the $\overline{\text{Fault}}$ pin.

The second lock-step configuration is called Cross-Coupled Checking. In general, one V_R4400MC is set to drive the data pins of one bus and the other the check pins (parity or ECC) of the same bus. The mode bits to use are 01 and 10. Both chips sample and compare what are on the busses and indicate any discrepancies through their $\overline{\text{Fault}}$ pins.

Additionally and for any configuration, the $\overline{\text{Fault}}$ pin also reports error conditions not covered by V_R4400MC exceptions. (These include input parity error at SysCmd, input addresses and data on SysAD for certain system interface operations.) Note that the fault detection mechanism related to this $\overline{\text{Fault}}$ pin does not cause any exceptions for the V_R4400MC. It continues to run regardless of the state of the pin. External logic must act on the fault and take appropriate action to interrupt or reset the V_R4400MC's.

8

8.1 Connections

For lock-step operations, most pins of a pair of V_R4400MC's are intended to be connected in parallel. The signal groups to be connected in parallel are:

- System Interface
- Secondary Cache Interface
- Interrupt Interface

Some signals in the following groups are NOT to be connected in parallel:

- Initialization Interface
 - ModeClock, ModeIn and $\overline{\text{Reset}}$
- JTAG Interface
 - JTDO, JTMS

The remaining signals in these groups can be connected either way, independently or in parallel, according to other considerations.

Finally, all of the signals in the Clock/Control Interface are NOT to be connected in parallel, except obviously V_{oo}P and G_{nd}P.

8.2 Modes of Operation

The Master/Checker mode of operation for each Vn4400MC is set via boot-time mode bits: System Interface Master (SIMasterMd) and Secondary Cache Master (SCMasterMd).

Table 8-1 Setting of SIMaster and SCMaster

SIMasterMd	SCMasterMd	Modes
0	0	Complete Master (required for single chip operation)
0	1	System Interface Master (SIMaster)
1	0	Secondary Cache Master (SCMaster, to pair with SIMaster)
1	1	Complete Listener (to pair with Complete Master)

Fig. 8-1 Master-Listener Configuration





Fig. 8-2 Cross-Coupled Checking Configuration



- (1) Signals connected in parallel and driven from the SIMaster are:
SysAD(63:0), SysCmd(8:0) and SCAPar(2:0)
- (2) Signals connected in parallel and driven from the SCmaster are:
SysADC(7:0), SysCmdP, ValidOut, Release, SCAddr(17:1), SCAddr0W,X,Y,Z,
SCOE, SCWrW,X,Y,Z, SCData(127:0), SCDChk(15:0),
SCTag(24:0), SCTChk(6:0), SCDCS, SCTCS

8.3 Fault Detection

Output miscomparison fault detection occurs at the end of bus cycles. The length of these cycles are determined by the appropriate boot-time mode bits.

At the System Interface, outputs are checked at the end of every System Interface cycle whenever the V_R4400MC System Interface is in master state.

At the Secondary Cache Interface, outputs are checked at the end of every read or write cycle.

A special note about the SCAPar(2:0) signals. They are delayed from the rest of the Secondary Cache Interface by exactly one PClock. This includes when they transition, and when they are checked. The transitions follow exactly one PClock after SCAddr transitions, and when the V_R4400MC internally determines it is changing between a read and a write cycle without an address change. They do not follow the timing of the SCWr signals, which are separately settable via the boot-time mode bits.

8.4 Fault Reporting

The external fault indication is reported through the $\overline{\text{Fault}}$ pin. This is a non-persistent signal which operates on the same bus cycles as the System Interface. That is, it transitions at System Interface boundaries, as determined by the PClock-to-SClock divisor from the boot-time mode bits.

There is an internal fault detection latency of 4 PClocks. Then the fault signal must synchronize with the System Interface. Any output faults detected and fully-propagated through the internal fault logic within the previous System Interface cycle will be reported in the current cycle.

Note that in the Complete Master mode, output fault reporting is disabled for Secondary Cache Interface pins, but enabled for System Interface pins (SysCmd, SysCmdP, SysAD, SysADC, $\overline{\text{ValidOut}}$ and $\overline{\text{Release}}$.) For System Interface pins, output faults are reported by comparing internal data.

8.5 Reset Operations

If the Master/Checker mode of the V_R4400MC is Complete Listener, SIMaster or SCMaster, an assertion of $\overline{\text{Reset}}$ is significant. On boot-up, the first time $\overline{\text{Reset}}$ is de-asserted, the Master/Checker mode is determined by the boot-time mode bits. On the next assertion of $\overline{\text{Reset}}$, the mode is forced to be a Forced Complete Master mode. Hence, if $\overline{\text{Reset}}$ is de-asserted, the V_R4400MC will be driving all outputs, regardless of the mode specified by the mode bits. This mode is maintained until $\overline{\text{Reset}}$ signal is asserted next time. If the $\overline{\text{Reset}}$ signal is asserted again, the V_R4400MC is set in the Master/Checker mode set by BTMC. In this way, the mode is changed between Master/Checker mode and Forced Complete Master mode each time $\overline{\text{Reset}}$ is asserted.

There is a slight difference between this Forced Complete Master mode and the normal Complete Master mode. That is, the $\overline{\text{Fault}}$ pin continues to report all output faults when the V_R4400MC was in the former mode, not just those of the System Interface.

8.6 Fault History Mechanism

Since both output faults and certain input faults are reported through the $\overline{\text{Fault}}$ pin, there are two fault history bits that record which one or both of them has occurred. These bits are cleared at every de-asserting transition of $\overline{\text{Reset}}$.

These bits are readable only when the $\overline{\text{Reset}}$ is asserted. The $\overline{\text{Fault}}$ pin changes from reporting live faults to indicating which fault history bit was set in the previous time period when $\overline{\text{Reset}}$ was de-asserted. The Modeln pin acts as the selector. If Modeln is 0, the $\overline{\text{Fault}}$ shows the state of the Output Fault History bit (inverted). If Modeln is 1, it shows the state of the Input Fault History bit (inverted).

There is also a mechanism to reset the fault history bits while the Vn4400MC is running. If the Modeln pin is asserted, these bits are cleared. This also means that for the duration when latching fault history bits are desired, Modeln must be maintained at 0.

The following table summarizes some of the above information.

Table 8-2 Fault History Mechanism Summary

Boot/Reset Controls	Modeln Pin	Fault History Bits	$\overline{\text{Fault}}$ Pin	Master/Checker Mode
VDDOk just asserted (0 $\rightarrow$ 1)	used as mode bits scan data	—	—	—
$\overline{\text{Reset}}$ just de-asserted (0 $\rightarrow$ 1)	—	cleared	—	—
$\overline{\text{Reset}}$ de-asserted (normal operation)	0	set and latch if fault	live faults reported	—
$\overline{\text{Reset}}$ de-asserted (normal operation)	1	cleared	live faults reported	—
$\overline{\text{Reset}}$ just asserted (1 $\rightarrow$ 0)	—	—	—	changed, toggling between mode bits and forced master
$\overline{\text{Reset}}$ just asserted (Vn4400MC in reset)	0	Output Fault History bit connected to $\overline{\text{Fault}}$ pin	—	—
$\overline{\text{Reset}}$ just asserted (Vn4400MC in reset)	1	Input Fault History bit connected to $\overline{\text{Fault}}$ pin	—	—

Remark Output fault : Contradiction of output data in Master/Checker mode
 Input fault : Data parity error of input data, ECC error

CHAPTER 9 PROCESSOR INTERRUPTS

The V_R4400MC processor supports six hardware interrupts, two software interrupts, and a non-maskable interrupt as described in V_R4000, V_R4400 USER'S MANUAL—ARCHITECTURE. Each hardware interrupt occurs when an external write request or $\overline{\text{Int0}}$ signal is input. The non-maskable interrupt is accessible via external write requests and a dedicated pin.

External writes to the processor are directed based on a processor internal address map to various processor internal resources. An external write to any address with SysAD[6..4] = 0 will write to the interrupt register. During the data cycle, SysAD[22..16] are the write enables for the 7 interrupt request bits (one non-maskable interrupt + 6 regular interrupts), and SysAD[6..0] are the values to be written into these bits. This allows us to set and clear any subset of the interrupt register with a single write request. The bit 0 of the interrupt register is bitwise ORed with the current value of the interrupt pin $\overline{\text{Int0}}$ and the result is directly readable as bit 10 of the Cause register.

When an interrupt is enabled, CPU can be altered, the interrupt exception is issued. The external interrupt corresponding to IP7 is alternated with a timer interrupt, and selected through BTMC interface.

NMI exception is issued by setting of bit6 of the interrupt register or asserting of the $\overline{\text{NMI}}$ pin.


The bit 15:10(IP7-IP3) of the Cause register are not cleared automatically after starting the handling of exception. Clear these bits as follows before the next interrupt will be enabled:

- An Interrupt request by the $\overline{\text{Int0}}$ pin: $\overline{\text{Int0}}$ is level-triggered. De-assert this pin.
- An Interrupt request by the external write request: Issue this request again and clear a corresponding bit of the interrupt register.

9

The interrupt request is generated through the logic shown below.

Fig. 9-1 Hardware Interrupt and NMI Request Generation Logic



The software interrupts are issued by setting the bits IP1 and IP0 of the Cause register to 1. The bits IM1 and IM0 of the Status register can mask these interrupts.

Keep issuing maskable interrupt requests until a jump to the exception vector. In an interrupt request by asserting the $\overline{\text{Int0}}$ pin, keep inputting low-level to the $\overline{\text{Int0}}$ pin. After the jump to the exception vector and starting the interrupt handling, clear the exception request before returning to the normal routine.

During the non-maskable interrupt, the exception request must not be cleared. The $\overline{\text{NMI}}$ pin is edge-triggered, so a single cycle assertion of this pin sets a non-maskable interrupt request.

Fig. 9-2 Interrupt Request Masking Logic




Table 9-1 Interrupt Requests Summary

Factor	Pins	External write requests		Interrupt register	Status register (IM bit)	Cause register (IP bit)
		request	mask			
NMI	$\overline{\text{NMI}}$	SysAD6	SysAD22	—	—	—
INT0	$\overline{\text{Int0}}$	SysAD0	SysAD16	bit10	bit10	bit10
INT1	—	SysAD1	SysAD17	bit11	bit11	bit11
INT2	—	SysAD2	SysAD18	bit12	bit12	bit12
INT3	—	SysAD3	SysAD19	bit13	bit13	bit13
INT4	—	SysAD4	SysAD20	bit14	bit14	bit14
INT5	—	SysAD5	SysAD21	bit15	bit15	bit15
SW0	—	—	—	bit8	bit8	bit8
SW1	—	—	—	bit9	bit9	bit9
Status		0: clear 1: request	0: disable 1: enable	0: clear 1: request	0: disable 1: enable	0: no request 1: pending

Remark See V4400, V4400 USER'S MANUAL—ARCHITECTURE for further details.

CHAPTER 10 PROCESSOR STATUS OUTPUTS

The V_R4400MC processor provides eight status outputs, Status(7:0), that change with each rising edge of MasterClock to indicate the processor's internal state during each of the two most recent PCycles. Status(7:0) is treated as two fields, Status(3:0) indicates the processor's internal state during the most recent PCycle and Status(7:4) indicates the processor's internal state during the PCycle preceding the most recent PCycle. The encoding of processor internal state for Status(7:4) or Status(3:0) is shown in Table 10-1 Encoding of Processor Internal State for Status(7:4) or Status(3:0). The four bit decode describes the instruction occupying the WB stage during a given PCycle.

Table 10-1 Encoding of Processor Internal State for Status(7:4) or Status(3:0)

<u>Status(7:4) or Status(3:0)</u>	<u>Processor internal state</u>	
0	Run cycle:	Other integer instruction
1	Run cycle:	Integer Load
2	Run cycle:	Integer Untaken Branch
3	Run cycle:	Integer Taken Branch
4	Run cycle:	Integer Store
5	Reserved	
6	Reserved	
7	Run cycle:	Killed by integer slip
8	Stall cycle:	Other stall type
9	Stall cycle:	Primary Instruction Cache
a	Stall cycle:	Primary Data Cache
b	Stall cycle:	Secondary Cache
c	Run cycle:	Floating Point instruction (except for load, store, conditional branch)
d	Run cycle:	Killed by branch, jump, ERET
e	Run cycle:	Killed by exception
f	Run cycle:	Killed by floating point slip

10

CHAPTER 11 CLOCKING

The V_A4400MC processor bases all internal and external clocking on the single clock input MasterClock. The processor generates the clock output SyncOut at the same frequency as MasterClock and aligns SyncIn with MasterClock. SyncOut must be connected to the clock input SyncIn so that the processor can compensate for output driver delays and input buffer delays in aligning SyncIn with MasterClock. The processor generates the clock output MasterOut at the same frequency as MasterClock and aligns MasterOut with SyncOut. MasterOut is provided for use in clocking external logic that must cycle at MasterClock frequency.

The processor generates the internal clock PClock for all internal latches and registers at twice the frequency of MasterClock and precisely aligns every other rising edge of PClock with the rising edge of MasterClock.

The processor divides PClock by a programmable divisor, programmed via the boot time mode control interface, to generate the internal clock SClock. SClock is used by the processor to sample data at the system interface and to clock data into the processor's system interface output registers. The rising edges of SClock are aligned with rising edges of PClock.

Data provided to the processor must be setup a minimum of t_{bs} nano-seconds (ns) before the rising edge of SClock and held valid for a minimum of t_{bh} ns after the rising edge of SClock. This setup and hold time is required for data to propagate through the processor's input buffers and meet the setup and hold times for the processor's input latches.

Data provided by the processor will become stable t_{oo} ns after the rising edge of SClock. This drive off time is the sum of the maximum delay through the processor's output drivers and the maximum clock to Q delay of the processor's output registers.

Certain processor inputs, specifically $\overline{V_{DD}Ok}$, $\overline{ColdReset}$, and $\overline{Reset}$ are sampled based on MasterClock while certain processor outputs, specifically Status(7:0) are driven out based on MasterClock. The same setup, hold, and drive off parameters, t_{bs} , t_{bh} , and t_{oo} , will apply to these inputs and outputs but with respect to MasterClock instead of SClock.

The values of t_{bs} , t_{bh} , and t_{oo} of V_A4400MC processor are tabulated in the data sheet (to be issued).

The processor generates two output clocks, RClock and TClock, at exactly the same frequency as SClock to be used by an external agent to sample and drive data. RClock is a receive clock that can be used by an external agent to clock its input registers. TClock is a transmit clock that can be used by an external agent to clock its output registers, and as the global system clock for the logic that makes up the external agent.

Although the frequencies of TClock and RClock are the same, RClock always precede TClock 1/4 cycle. The relationship between RClock and TClock is independent of the delay between SyncIn and SyncOut.

The frequencies of TClock and SClock are the same, and their relation with SyncIn and SyncOut is as follows:

When SyncIn and SyncOut are connected:


The edge of TClock and edge of SClock accurately correspond to each other.

If there is a delay between SyncIn and SyncOut:

TClock seems to change faster than SClock to the external agent.

This delay also affects MasterClock. If the delay between SyncIn and SyncOut is the same as the delay between the processor and external agent, the last TClock of the external agent corresponds to SClock and MasterClock.

Fig. 11-1 Clocks and PLL Network



Caution "PLL-Off mode" is just for testing. NEC does not guarantee functionality of this mode.

The alignment of SyncOut, PClock, SClock, TClock, and RClock is accomplished by the processor with internal Phase Locked Loop (PLL) circuits that generate aligned clocks. A functional block diagram of the internal clocks and PLL network is shown in Fig.10-1. PLL circuits by their nature are only capable of generating aligned clocks for MasterClock frequencies in a limited range. Minimum and maximum frequencies for MasterClock for various speed ratings of the V_R4400MC processor are tabulated in the data sheet (to be issued). Clocks generated using PLL circuits contain some inherent inaccuracy in their alignment with the MasterClock called jitter. That is, a clock aligned with MasterClock by the processor's PLL circuits may lead or trail MasterClock by some maximum amount referred to as the maximum jitter. Maximum jitter for the clocks generated by various speed ratings of the V_R4400MC processor is tabulated in the data sheet (to be issued). For optimum operation, the traces and loads on SyncOut, MasterOut, TClock, and RClock should be matched.

The PClock to SClock Divisor can be selected out of 1/2, 1/3, 1/4, 1/6, and 1/8 in the V_R4400MC through the BTMC interface on reset.

The relationship of MasterClock, SyncOut, PClock, SClock, TClock, and RClock, and the characteristics of data on the SysAD bus when data is driven by the processor and received by the processor is illustrated in Fig. 11-2 Processor Clocks, PClock to SClock Divisor of 2 and frequency of PClock divided by two while in Fig. 11-3 Processor Clocks, PClock to SClock Divisor of 4, SClock, TClock, and RClock are programmed to the frequency of PClock divided by four.

Fig. 11-2 Processor Clocks, PClock to SClock Divisor of 2





Fig. 11-3 Processor Clocks, PClock to SClock Divisor of 4




11.1 Clock Interfacing to a Phase Locked System

When the Vr4400MC processor is employed in a phase locked system all of the components of the system must phase lock their operation to a common MasterClock. In such a system the delivery of data and sampling of data will have common characteristics for all components with perhaps different delay values for the components. The **transmission time**, the amount of time a signal has to propagate along the trace from one component to another, between any two components A and B of a phase locked system can be calculated from the following equation:

$$\begin{aligned} \text{Transmission Time} = & (\text{SClock period}) - (t_{bo} \text{ for A}) - (t_{bs} \text{ for B}) \\ & - (\text{Clock Jitter for A Max}) - (\text{Clock Jitter for B Max}) \end{aligned}$$

A block level diagram of a phase locked system employing the Vr4400MC processor is shown in Fig. 11-4 **Phase Locked System Employing the Vr4400MC Processor**.

Fig. 11-4 Phase Locked System Employing the Vr4400MC Processor



11.2 Clock Interfacing to a System Without Phase Lock

When the Vr4400MC processor is employed in a system in which the other components are not capable of phase lock to a single MasterClock the output clocks RClock and TClock may be used to clock the remainder of the system. Two clocking methodologies are possible using RClock and TClock: the first better suited to communication with an external agent built from gate arrays and the second better suited to communication with an external agent built from discrete CMOS logic devices.

In the first clocking methodology, tailored for communication with an external agent built from gate arrays both RClock and TClock are used for clocking within the gate arrays. RClock is provided specifically so that a gate array may buffer it internally and use the buffered version to clock registers that sample Vr4400MC outputs. These sample registers should be immediately followed by staging registers that are clocked by an internally buffered version of TClock. The buffered version of TClock should be used as the global system clock for the logic inside the gate array and as the clock for all registers that drive Vr4400MC inputs.

Requiring staging registers following the registers that sample Vr4400MC outputs places a constraint on the delay of the input registers and the setup time of the synchronizing registers inside the gate arrays. The sum must be less than 25% of the RClock and TClock period minus the maximum clock jitter of both RClock and TClock minus the maximum delay mismatch for the internal clock buffers on RClock and TClock.

The transmission time for a signal from the Vr4400MC to an external agent composed of gate arrays in a system without phase lock can be calculated from the following equation:


$$\begin{aligned}
 \text{Transmission Time} = & (75\% \text{ of TClock period}) - (\text{tbo for Vr4400MC}) \\
 & + (\text{External Clock Buffer Delay Min}) \\
 & - (\text{External Sample Register Setup Time}) \\
 & - (\text{Clock Jitter for Vr4400MC Internal Clocks Max}) \\
 & - (\text{Clock Jitter for RClock Max})
 \end{aligned}$$

The transmission time for a signal from an external agent composed of gate arrays to the V_A4400MC in a system without phase lock can be calculated from the following equation:

$$\begin{aligned} \text{Transmission Time} = & (\text{TClock period}) - (\text{tos for V}_{A4400\text{MC}}) \\ & - (\text{External Clock Buffer Delay Max}) \\ & - (\text{External Output Register Clock Delay Max}) \\ & - (\text{Clock Jitter for TClock Max}) \\ & - (\text{Clock Jitter for V}_{A4400\text{MC}} \text{ Internal Clocks Max}) \end{aligned}$$

A block level diagram of a system without phase lock employing the V_A4400MC processor and an external agent composed of a gate array is shown in Fig. 11-5 System Without Phase Lock Employing the V_A4400MC Processor (a).

Fig. 11-5 System Without Phase Lock Employing the Vr4400MC Processor (a)



In the second clocking methodology tailored for communication with an external agent built from discrete CMOS logic devices, matched delay clock buffers are used to allow the V_A4400MC to generate aligned clocks for the external logic. One of the matched delay clock buffers is inserted in the processor's SyncOut SyncIn clock alignment path. This has the effect of skewing SyncOut, MasterOut, RClock, and TClock to lead MasterClock by the delay of the matched delay clock buffer while leaving PClock aligned with MasterClock. The remaining matched delay clock buffers can be used to generate a buffered version of TClock that will be aligned with MasterClock. The alignment error of the buffered version of TClock will be the sum of the maximum delay mismatch of the matched delay clock buffers and the maximum clock jitter of TClock. The buffered version of TClock will be used to clock registers that sample V_A4400MC outputs, as the global system clock for the discrete logic that forms the external agent, and to clock registers that drive V_A4400MC inputs.

The transmission time for a signal from the V_A4400MC to an external agent composed of discrete CMOS logic devices can be calculated from the following equation:

$$\begin{aligned} \text{Transmission Time} = & (\text{TClock period}) - (\text{t}_{\text{do}} \text{ for V}_{\text{A}}4400\text{MC}) \\ & - (\text{External Sample Register Setup Time}) \\ & - (\text{External Clock Buffer Delay Mismatch Max}) \\ & - (\text{Clock Jitter for V}_{\text{A}}4400\text{MC Internal Clocks Max}) \\ & - (\text{Clock Jitter for TClock Max}) \end{aligned}$$


The transmission time for a signal from an external agent composed of discrete CMOS logic devices can be calculated from the following equation:

$$\begin{aligned} \text{Transmission Time} = & (\text{TClock period}) - (\text{t}_{\text{ds}} \text{ for V}_{\text{A}}4400\text{MC}) \\ & - (\text{External Output Register Clock Delay Max}) \\ & - (\text{External Clock Buffer Delay Mismatch Max}) \\ & - (\text{Clock Jitter for V}_{\text{A}}4400\text{MC Internal Clocks Max}) \\ & - (\text{Clock Jitter for TClock Max}) \end{aligned}$$

Note that using this clocking methodology the hold time of data driven from the V_A4400MC to an external sampling register is a critical parameter. In order to guarantee hold time, the minimum output delay of the V_A4400MC must be greater than the sum of the minimum hold time for the external sampling register, the maximum clock jitter for V_A4400MC internal clocks, the maximum clock jitter for TClock, and the maximum delay mismatch of the external clock buffers.

A block level diagram of a system without phase lock employing the V_A4400MC processor and an external agent composed of both a gate array and discrete CMOS logic devices is shown in Fig. 11-6 **System Without Phase Lock Employing the V_A4400MC Processor (b)**.

Fig. 11-6 System Without Phase Lock Employing the Vn4400MC Processor (b)



CHAPTER 12 OUTPUT BUFFER di/dt CONTROL MECHANISM

The speed of the V_R4400MC output drivers is controlled by a negative feedback loop that insures that the drive off times are only as fast as necessary to meet the system requirement of single cycle transfers. This guarantees the minimum ground bounce due to the $L \cdot di/dt$ of the switching buffers consistent with the system timing requirements. Four bits are used to control each of the pull-up and pull-down delays. They are initially set to the values in the mode bits InitN[3..0] for pull-up and InitP[3..0] for pull-down. When the di/dt control mechanism is enabled, InitN and InitP should be set to their slowest values.

Under normal conditions, the di/dt control mechanism is expected to be constantly enabled so that it can compensate the output buffer delay for any changes in the temperature or power supply voltage. The EnbIDPLL mode bit should be set for this mode of operation.

For situations where the jitter associated with the operation of the di/dt control mechanism cannot be tolerated and where the variation in temperature and supply voltage after ColdReset is expected to be small, the di/dt control mechanism can be instructed to lock only during ColdReset and thereafter retain its control values. The EnbIDPLL mode bit should be set and the EnbIDPLL mode bit should be cleared for this mode of operation.

In addition, if both the EnbIDPLL and EnbIDPLL mode bits are cleared, the speed of the output buffers can be set with the InitP[3..0] and InitN[3..0] mode bits.

The drive off delays can be set through the mode bits. Currently, delays of 0.5T, 0.75T, and T are supported corresponding to the Drv0_50, Drv0_75, and Drv1_00 mode bits where T is the MasterClock period. For example, in the Drv0_75 mode, the entire signal transmission path including the clock-to-Q, output buffer drive time, board flight time, input buffer delay, and setup time will be traversed in $0.75 \cdot$ the MasterClock period plus or minus the jitter due to the di/dt control mechanism.

All output drivers on the V_R4400MC, with the exception of the clock drivers, are controlled by the di/dt control mechanism.


The V_R4400MC determines the worst case propagation delay from a V_R4400MC output driver to a receiving device by measuring the transmission line delay of the trace that connects the V_R4400MC IO_Out and IO_In pins. This representative trace must have one and a half times the length and approximately the same capacitive loading as the worst case trace on any V_R4400MC output.

The designer determines the trace characteristics by:

- measuring the longest path from a V_R4400MC output driver to a receiving device: L
- calculating the maximum capacitive loading on any signal pin: C
- connecting an "incident wave" trace of length L with a capacitive loading of C between the IO_In and IO_Out pins of the V_R4400MC
- and connecting a "reflected wave" trace of length L/2 to the IO_In pin of the V_R4400MC.

A V_R4400MC with appropriate traces connected to the IO_In and IO_Out pins is illustrated in Fig. 12-1 IO_In/IO_Out Board Trace.

Fig. 12-1 IO_In/IO_Out Board Trace




CHAPTER 13 PLL PASSIVE COMPONENTS

The Phase Locked Loops require several passive components for proper operation. These passive components are attached to the PLLCap0, PLLCap1, V_{DD}P, and GndP pins and are illustrated below. The capacitors for the PLLCap0 and PLLCap1 pads are connected to either GndP or V_{DD}P. Note that C2 and both Cp capacitors are incorporated into package designs as surface mounted chip capacitors, illustrated in Fig.12-1. Note that the capacitors connecting the PLLCap0 and PLLCap1 pins to either V_{DD}P or GndP are to be provided externally.


It is essential to isolate the analog power and ground (V_{DD}P/GndP) from the other power and ground pins. Initial evaluations with R=5 Ω , C1=1 nF, C2=0.1 μ F, C3=10 μ F, and Cp=470 pF have yielded good results on test boards. The inductors (L) may be excluded, since noise filtering is application specific. The inductors may improve noise reduction in some applications. The suggested value for the inductor is 1000 μ H. Since the optimum values for the filter components depend on the application and the system noise environment, these values should be considered as starting points for further experimentation within the application specific context. In addition, the chokes (inductors) can be considered as an alternative to the resistors for power supply filtering.

Fig. 13-1 PLL Passive Components



- Remarks
1. ----->: Unconnected in the present Vt4400MC.
 2. In the present Vt4400MC, PLLCap0 is connected to GndP, and PLLCap1 to VDDP.

Fig. 13-2 Top View of 447-pin PGA Package



x: Gnd - VDD Bypass Caps
 C2: GndP - VDDP Bypass Caps
 %1, %2: PLL Caps

CHAPTER 14 JTAG INTERFACE

The V_A4400MC processor provides a boundary scan interface using the industry standard JTAG protocol.

14.1 JTAG Interface Signal Summary

- JTDI: (i) JTAG serial data in.
JTDO: (o) JTAG serial data out.
JTMS: (i) JTAG command signal.
JTCK: (i) JTAG serial clock input.

14.2 JTAG Functionality

The JTAG boundary scan mechanism is intended to provide a capability for testing the interconnect between the V_A4400MC processor, the printed circuit board to which it is attached, and the other components on the board. In addition, the JTAG boundary scan mechanism is intended to provide a rudimentary capability for low speed logical testing of the secondary cache RAMs. The JTAG boundary scan mechanism is not intended to provide any capability for testing the V_A4400MC processor itself.

In accordance with the JTAG specification the V_A4400MC processor contains a TAP controller, JTAG instruction register, JTAG boundary scan register, and JTAG bypass register. However, the V_A4400MC JTAG implementation provides only the external test functionality of the boundary scan register.

14.2.1 JTAG Test Access Port (TAP)

The JTAG Test Access Port consists of the 4 pins described above. Data is serially scanned into one of the three registers (Instruction register, Bypass register, Boundary Scan register) from the JTDI pin, and is scanned out from the selected one of these registers onto the JTDO pin. The JTDI input feeds the LSB of the selected register, and the MSB of the selected register appears on the JTDO output. The JTMS input controls the state transitions of the main TAP controller state machine.

Data on the JTDI and JTMS pins is sampled on the rising edge of the JTCK input clock signal. Data on the JTDO pin changes on the falling edge of the JTCK clock signal.

14.2.2 JTAG TAP Controller

The V_A4400MC implements the 16-state JTAG TAP controller as defined in the IEEE JTAG specification.

The TAP controller state machine can be put in its Reset state in one of two ways. Deassertion of the V_{DD}Ok input will reset the TAP controller. Keeping the JTMS input signal asserted through five consecutive rising edges of the JTCK clock input will also send the TAP controller state machine into its Reset state. In either case, keeping JTMS asserted will maintain the Reset state.

14.2.3 Instruction Register

The V_A4400MC's JTAG instruction register is three bits wide and is encoded as follows.

MSB...LSB	Selected Data Register
0 0 0	Boundary Scan register (External Test only)
0 0 1	Bypass Register
0 1 x	Bypass Register
1 x x	Bypass Register

The Instruction register is composed of two stages – the shift register stage and the parallel output latch. When the TAP controller is in the Reset state, the value 7 (111) is loaded into the parallel output latch, thus selecting the Bypass register as the default. When the TAP controller is in the Capture-IR state, the value 4 (100) is loaded into the shift register stage. When the TAP controller is in the Shift-IR state, data is serially shifted into the shift register stage of the Instruction register from the JTDI input pin, and the MSB of the Instruction register's shift register stage is shifted out onto the JTDO pin. When the TAP controller is in the Update-IR state, the current data in the shift register stage is loaded into the parallel output latch.

14.2.4 Bypass Register

The Bypass Register is one bit wide. When the TAP controller is in the Shift-DR (Bypass) state, the data on the JTDI pin is shifted into the bypass register, and the bypass register's output is shifted out onto the JTDO output pin.

14.2.5 Boundary Scan Register

The Boundary Scan register is 319 bits wide. The three most significant bits control the output enables on the various bidirectional buses. The most significant bit is the JTAG output enable bit for the SysAD, SysADC, SysCmd and SysCmdP buses. The next most significant bit is the JTAG output enable for the SCDATA and SCDCHK buses. The third most significant bit is the JTAG output enable for the SCTAG and SCTCHK buses. The remaining 316 bits correspond to 316 signal pads of the V_{R4400MC}. The scan order of these bits is listed in **APPENDIX C JTAG ORDERING** at the end of this document.

When the TAP controller is in the Reset state, the three most significant bits of the Boundary Scan register are set to "0" (the default JTAG output enable control on all the bidirectional pins is to disable the outputs). When the TAP controller is in the Capture-DR (Boundary Scan) state, the data currently present on all the V_{R4400MC}'s input and I/O pins are latched into the Boundary Scan register. The Boundary Scan register bits corresponding to output pins are arbitrary in this state and must not be checked during the scan out process. When the TAP controller is in the Shift-DR (Boundary Scan) state, data is serially shifted into the Boundary Scan register from the JTDI pin, and the contents of the Boundary Scan register are shifted out onto the JTDO pin. When the TAP controller is in the Update-DR (Boundary Scan) state, the current data in the Boundary Scan register is latched into its parallel output latch, and the bits corresponding to output pins and those IO pins whose outputs are enabled (by the three MSBs of the Boundary Scan register) are enabled onto the V_{R4400MC}'s pins.

14.3 Implementation Specific Details

- The MasterClock, MasterOut, SyncIn and SyncOut pads do not have JTAG.
- Some pairs of output pads share a single JTAG bit. These are: SCAAddr0W and SCAAddr0X, SCAAddr0Y and SCAAddr0Z, $\overline{SCWrW}$ and $\overline{SCWrX}$, $\overline{SCWrY}$ and $\overline{SCWrZ}$, TClock[0] and TClock[1], RClock[0] and RClock[1].
- All input pads data are first latched into a Processor Clock based register in the pad cell before they are captured into the Boundary Scan register in the Capture-DR (Boundary Scan) state. When the Phase locked loop is disabled, the processor clock is half the frequency of MasterClock. Therefore the data setup required at the input pads is greater than two MasterClock periods before the rising edge of JTCK when the TAP controller is in the Capture-DR (Boundary Scan) state.
- The output enable controls generated from the three most significant bits of the Boundary Scan register are latched into a Processor Clock based register before they actually enable the data onto the pads. Therefore the delay from the rising edge of JTCK in the Update-DR (Boundary Scan) state to data valid at the output pins of the chip is greater than two MasterClock periods.

CHAPTER 15 PIN SUMMARY

Secondary cache interface pins:

- SCData(127:0):** (i/o) A 128-bit bus used to read or write cache data from/to the secondary cache.
- SCDChk(15:0):** (i/o) A 16-bit bus which conveys two ECC fields that cover the upper or lower 64 bits of the SCData from/to the secondary cache.
- SCTag(24:0):** (i/o) A 25-bit bus used to read or write cache tags from/to the secondary cache.
- SCTChk(6:0):** (i/o) A 7-bit bus which conveys an ECC field that covers the SCTag from/to the secondary cache.
- SCAddr(17:1):** (o) A 17-bit bus which addresses the secondary cache.
- SCAddr0Z:** (o) Bit 0 of the secondary cache address.
- SCAddr0Y:** (o) Bit 0 of the secondary cache address.
- SCAddr0X:** (o) Bit 0 of the secondary cache address.
- SCAddr0W:** (o) Bit 0 of the secondary cache address.
- SCAPar(2:0):** (o) The secondary cache address even parity bus cover the following bits:
- SCAPar(2) 7 bits: $\overline{\text{SCWr}}$, SCAddr(17:12)
 - SCAPar(1) 7 bits: $\overline{\text{SCDCS}}$, SCAddr(11:6)
 - SCAPar(0) 7 bits: $\overline{\text{SCTCS}}$, SCAddr(5:0)
- $\overline{\text{SCOE}}$:** (o) A signal which enables the outputs of the secondary cache RAMs.
- $\overline{\text{SCWrZ}}$:** (o) Secondary cache write enable.
- $\overline{\text{SCWrY}}$:** (o) Secondary cache write enable.
- $\overline{\text{SCWrX}}$:** (o) Secondary cache write enable.
- $\overline{\text{SCWrW}}$:** (o) Secondary cache write enable.
- $\overline{\text{SCDCS}}$:** (o) A signal which enables the chip select pins of the secondary cache RAMs associated with SCData and SCDChk.
- $\overline{\text{SCTCS}}$:** (o) A signal which enables the chip select pins of the secondary cache RAMs associated with SCTag and SCTChk.

System interface pins:

- SysAD(63:0):** (i/o) A 64-bit bus used for address and data transmission between the processor and an external agent.
- SysADC(7:0):** (i/o) An 8-bit bus containing check bits for the SysAD bus.
- SysCmd(8:0):** (i/o) A 9-bit bus used for command and data identifier transmission between the processor and an external agent.
- SysCmdP:** (i/o) A single even parity bit over the SysCmd bus. When the system interface is set in the parity mode, the processor indicates the ECC error of the secondary cache by using this bit.
- $\overline{\text{ValidIn}}$:** (i) Signals that an external agent is driving a valid address or valid data on the SysAD bus and a valid command or data identifier on the SysCmd bus during this cycle.
- $\overline{\text{ValidOut}}$:** (o) Signals that the processor is driving a valid address or valid data on the SysAD bus and a valid command or data identifier on the SysCmd bus during this cycle.
- $\overline{\text{ExtRqst}}$:** (i) Signals that the system interface needs to submit an external request.

- $\overline{\text{Release}}$:** (o) Signals that the processor is releasing the system interface to slave state.
- $\overline{\text{RdRdy}}$:** (i) Signals that an external agent is capable of accepting a processor read, invalidate, or update request in both non-overlap and overlap mode or a read followed by a potential invalidate or update request in overlap mode.
- $\overline{\text{WrRdy}}$:** (i) Signals that an external agent is capable of accepting a processor write request in both non-overlap and overlap mode.
- $\overline{\text{IvdAck}}$:** (i) Signals that a processor invalidate or update request has completed successfully.
- $\overline{\text{IvdErr}}$:** (i) Signals that a processor invalidate or update request has completed unsuccessfully.

Interrupt pin:

- $\overline{\text{Int}}(0)$:** (i) One of six general processor interrupts, bit-wise ORed with bit 0 of the interrupt register.

Non-maskable interrupt pin:

- $\overline{\text{NMI}}$:** (i) Non-maskable interrupt, ORed with bit 6 of the interrupt register.

Boot time mode control interface pins:

- ModeIn:** (i) Serial boot mode data in.
- ModeClock:** (o) Serial boot mode data clock out at the system clock frequency divided by 256.

JTAG interface pins:

- JTDI:** (i) JTAG serial data in.
- JTDO:** (o) JTAG serial data out.
- JTMS:** (i) JTAG command signal, signals that the serial data in is command data.
- JTCK:** (i) JTAG serial clock input. Make this signal low when the JTAG interface is not used.

Maintenance pins:

- TClock(1:0):** (o) Two identical transmit clocks at the operation frequency of the system interface.
- RClock(1:0):** (o) Two identical receive clocks at the operation frequency of the system interface.
- MasterClock:** (i) Master clock input at the operation frequency of the processor.
- MasterOut:** (o) Master clock output aligned with MasterClock.
- SyncOut:** (o) Synchronization clock output.
- SyncIn:** (i) Synchronization clock input.
- IOOut:** (o) Output slew rate control feedback loop output. Must be connected to IOIn through a delay loop that models the IO path from the V_{R4400MC} to an external agent.
- IOIn:** (i) Output slew rate control feedback loop input.
- V_{DD}Ok:** (i) When asserted, this signal indicates to the V_{R4400MC} that the power supply voltage has been within the specified range for more than 100 milliseconds and will remain stable. The assertion of V_{DD}Ok will initiate the reading of the boot time mode control serial stream.
- $\overline{\text{ColdReset}}$:** (i) This signal must be asserted for a power on reset or a cold reset. The clocks SClock, TClock, and RClock begin to cycle and are synchronized with the deassertion edge of $\overline{\text{ColdReset}}$. $\overline{\text{ColdReset}}$ must be de-asserted synchronously with MasterClock.

Reset:	(i) This signal must be asserted for any reset sequence. It may be asserted synchronously or asynchronously for a cold reset, or synchronously to initiate a warm reset. <u>Reset</u> must be de-asserted synchronously with MasterClock.
Fault:	(o) Mismatch output of boundary comparators.
V_{DD}P:	(i) Quiet V _{DD} for the internal phase locked loop.
GndP:	(i) Quiet Gnd for the internal phase locked loop.
Status(7:0):	(o) An 8-bit bus that indicates the current operation status of the processor.
V_{DD}Sense:	(i/o) This is a special pin used only in component testing and characterization. It provides a separate, direct connection from the on-chip V _{DD} node to a package pin without attaching to the in-package power planes. Test fixtures treat V _{DD} Sense as an analog output pin; the voltage at this pin directly shows the behavior of the on-chip V _{DD} . Thus, characterization engineers can easily observe the effects of di/dt noise, transmission line reflections, etc. V _{DD} Sense should be connected to V _{DD} in functional system designs.
GndSense:	(i/o) GndSense provides a separate, direct connection from the on-chip Gnd node to a package pin without attaching to the in-package ground planes. GndSense should be connected to Gnd in functional system designs.

APPENDIX A SUB-BLOCK ORDERING

Sub-block ordering is an order for transmitting the data elements that form block of data when the data element transmitted first is not the data element at the beginning of the block. Sub-block ordering causes the data elements of the block to be transmitted in an order that fills out sub-blocks of increasing size. For the V44400MC, the smallest data element of a block transfer is a double word, therefore, the double word at the target address is transferred first, followed by the double word that fills out the quad word that contains the starting double word. Next the quad word that fills out an octal word containing the starting quad word is transferred in the same order as the first quad word, followed by the octal word that fills out a hex word containing the starting octal word in the same order as the first octal word, and so on through sub-blocks of increasing size until the entire block has been transferred.

Perhaps an easier way to consider sub-block ordering is to look at a method for generating the addresses, within the block, of the double words to be transferred for sub-block ordering. A simple method for generating such addresses is to bit-wise XOr the starting double word address with the output of a binary counter that is counting the double words in the block starting at double word zero.

The following tables illustrate the sequence of double words transferred using sub-block ordering for a thirty-two word block based on three different starting block addresses. For these illustrations the double words in the block will be identified by their block addresses. The block address for each double word in a block is derived by numbering the double words in the block sequentially starting with zero.

The tables also include a binary count of the double words in the block to illustrate the XOr relationship between this count, the starting address and the block addresses of the double words transferred.

Table A-1 Sequence of Double Words Transferred Using Sub-block Ordering (a)

<u>Cycle</u>	<u>Starting block address</u>	<u>Binary count</u>	<u>Double word transferred</u>
1	0010	0000	0010
2	0010	0001	0011
3	0010	0010	0000
4	0010	0011	0001
5	0010	0100	0110
6	0010	0101	0111
7	0010	0110	0100
8	0010	0111	0101
9	0010	1000	1010
10	0010	1001	1011
11	0010	1010	1000
12	0010	1011	1001
13	0010	1100	1110
14	0010	1101	1111
15	0010	1110	1100
16	0010	1111	1101

A

Table A-2 Sequence of Double Words Transferred Using Sub-block Ordering (b)

<u>Cycle</u>	<u>Starting block address</u>	<u>Binary count</u>	<u>Double word transferred</u>
1	1011	0000	1011
2	1011	0001	1010
3	1011	0010	1001
4	1011	0011	1000
5	1011	0100	1111
6	1011	0101	1110
7	1011	0110	1101
8	1011	0111	1100
9	1011	1000	0011
10	1011	1001	0010
11	1011	1010	0001
12	1011	1011	0000
13	1011	1100	0111
14	1011	1101	0110
15	1011	1110	0101
16	1011	1111	0100

Table A-3 Sequence of Double Words Transferred Using Sub-block Ordering (c)

<u>Cycle</u>	<u>Starting block address</u>	<u>Binary count</u>	<u>Double word transferred</u>
1	0101	0000	0101
2	0101	0001	0100
3	0101	0010	0111
4	0101	0011	0110
5	0101	0100	0001
6	0101	0101	0000
7	0101	0110	0011
8	0101	0111	0010
9	0101	1000	1101
10	0101	1001	1100
11	0101	1010	1111
12	0101	1011	1110
13	0101	1100	1001
14	0101	1101	1000
15	0101	1110	1011
16	0101	1111	1010

APPENDIX B EVEN PARITY

The descriptions of parity as even or odd often have different meanings to different people. Even parity is defined for the V_R4400MC as follows:

For a field of n bits protected by a single parity bit, if the number of ones among the n data bits is an even number, then the even parity bit will be a 0. If the number of ones among the n data bits is an odd number, then the even parity bit will be a 1. For example, if all n of the data bits are 0, the parity bit will also be a 0 if there are no data bits in error. If all n of the data bits are 1, and the number of data bits n is an odd number, then the parity bit will be a 1 if there are no data bits in error.

APPENDIX C JTAG ORDERING

The following list is the order of the pins associated with the JTAG Boundary Scan Register starting from JTDI and ending at JTDO.

SCDChk[13]	SysAD[29]
SysADC[1]	SCData[125]
SCDChk[1]	<u>Reset</u>
SysADC[5]	SCTag[20]
SCDChk[5]	SCData[93]
Status[0]	SCData[60]
Status[1]	SysAD[60]
Status[2]	SCData[28]
Status[3]	SysAD[28]
<u>IvdErr</u>	SCData[124]
Status[4]	<u>ColdReset</u>
<u>IvdAck</u>	SCTag[21]
Status[5]	SCData[92]
Status[6]	SCData[59]
Status[7]	SysAD[59]
SCDChk[7]	SCData[27]
SysADC[7]	SysAD[27]
SCDChk[3]	SCData[123]
SysADC[3]	IOIn
SCDChk[15]	SCTag[22]
VddOk	SCData[91]
SCTag[16]	SCData[58]
SCDChk[11]	SysAD[58]
SCData[63]	SCData[26]
SysAD[63]	SysAD[26]
SCData[31]	SCData[122]
SysAD[31]	IOOut
SCData[127]	SCTag[23]
SCTag[17]	SCData[90]
SCData[95]	SCData[57]
SCData[62]	SysAD[57]
SysAD[62]	SCData[25]
SCData[30]	SysAD[25]
SysAD[30]	SCData[121]
SCData[126]	<u>GrpRun</u>
SCTag[18]	SCTag[24]
SCData[94]	SCData[89]
RClock[1..0] (share the same JTAG bit)	SCData[56]
SCTag[19]	SysAD[56]
SCData[61]	SCData[24]
SysAD[61]	SysAD[24]
SCData[29]	SCData[120]

<u>GrpStall</u>	<u>ValidIn</u>
SCTChk[0]	SCTChk[5]
SCData[88]	SCData[83]
SCDChk[6]	SCAddr0W,X (share the same JTAG bit)
SysADC[6]	SCAddr0Y,Z (share the same JTAG bit)
SCDChk[2]	SCAddr[1]
SysADC[2]	SCData[50]
SCDChk[14]	SysAD[50]
<u>NMI</u>	SCData[18]
SCTChk[1]	SysAD[18]
SCDChk[10]	SCData[114]
SCData[55]	<u>Int[0]</u>
SysAD[55]	SCTChk[6]
SCData[23]	SCData[82]
SysAD[23]	SCData[49]
SCData[119]	SysAD[49]
<u>Release</u>	SCData[17]
SCTChk[2]	SysAD[17]
SCData[87]	SCData[113]
SCData[54]	SCAddr[2]/ <u>Int[1]</u>
SysAD[54]	SCAddr[3]
SysAD[22]	SCData[81]
<u>Modeln</u>	SCData[48]
SCData[22]	SysAD[48]
<u>RdRdy</u>	SCData[16]
SCData[118]	SysAD[16]
SCData[86]	SCData[112]
SCData[53]	SCAddr[4]/ <u>Int[2]</u>
SysAD[53]	SCAddr[5]
SCData[21]	SCData[80]
SysAD[21]	SCAddr[6]
SCData[117]	SCAddr[7]
<u>ExtRqst</u>	SCAddr[8]
SCTChk[3]	SCAddr[9]
SCData[85]	SCAddr[10]
SCData[52]	SCAddr[11]
SysAD[52]	RFU
SCData[20]	SCAddr[12]
SysAD[20]	SCAddr[13]
SCData[116]	SCAddr[14]
<u>ValidOut</u>	SCAddr[15]
SCTChk[4]	SCAddr[16]
SCData[84]	SCAddr[17]
SCData[51]	SCData[64]
SysAD[51]	SCAPar[0]
SCData[19]	SCAPar[1]/ <u>Int[3]</u>
SysAD[19]	SCData[96]
SCData[115]	SysAD[0]

SCData[0]	ModeClk
SysAD[32]	SCData[102]
SCData[32]	SysAD[6]
SCData[65]	SCData[6]
SCAPar[2]	SysAD[38]
$\overline{SCO\bar{E}}/int[4]$	SCData[38]
SCData[97]	SCData[71]
SysAD[1]	SCTag[4]
SCData[1]	SysCmd[3]
SysAD[33]	SCData[103]
SCData[33]	SysAD[7]
SCData[66]	SCData[7]
$\overline{SCDCS}$	SysAD[39]
$\overline{SCTCS}/int[5]$	SCData[39]
SCData[98]	SCDChk[8]
SysAD[2]	SCTag[5]
SCData[2]	SysCmd[4]
SysAD[34]	SCDChk[12]
SCData[34]	SysADC[0]
SCTag[0]	SCDChk[0]
$\overline{SCWrW}, X$ (share the same JTAG bit)	SysADC[4]
$\overline{SCWrY}, Z$ (share the same JTAG bit)	SCDChk[4]
SCData[67]	SCData[72]
SCTag[1]	SCTag[6]
SysCmd[0]	SysCmd[5]
SCData[99]	SCData[104]
SysAD[3]	SysAD[8]
SCData[3]	SCData[8]
SysAD[35]	SysAD[40]
SCData[35]	SCData[40]
SCData[68]	SCData[73]
SCTag[2]	SCTag[7]
SysCmd[1]	SysCmd[6]
SCData[100]	SCData[105]
SysAD[4]	SysAD[9]
SCData[4]	SCData[9]
SysAD[36]	SysAD[41]
SCData[36]	SCData[41]
SCData[69]	SCData[74]
SCTag[3]	SCTag[8]
SysCmd[2]	SysCmd[7]
SCData[101]	SCData[106]
SysAD[5]	SysAD[10]
SCData[5]	SCData[10]
SysAD[37]	SysAD[42]
SCData[37]	SCData[42]
SCData[70]	SCData[75]
$\overline{WrRdy}$	SCTag[9]

SysCmd[8]
 SCDData[107]
 SysAD[11]
 SCDData[11]
 SysAD[43]
 SCDData[43]
 SCDData[76]
 SCTag[10]
 SysCmdP
 SCDData[108]
 SysAD[12]
 SCDData[12]
 SysAD[44]
 SCDData[44]
 SCDData[77]
 SCTag[11]
 Fault
 SCDData[109]
 SysAD[13]
 SCDData[13]
 SysAD[45]
 SCDData[45]
 SCTag[12]
 TClock[1..0] (share the same JTAG bit)
 SCDData[78]
 SCTag[13]
 SCDData[110]
 SysAD[14]
 SCDData[14]
 SysAD[46]
 SCDData[46]
 SCDData[79]
 SCTag[14]
 SCDData[111]
 SysAD[15]
 SCDData[15]
 SysAD[47]
 SCDData[47]
 SCDChk[9]
 SCTag[15]
 SCTag_OE (JTAG output enable control for SCTag and SCTChk buses)
 SCDData_OE (JTAG output enable control for SCDData and SCDChk buses)
 SysAD_OE (JTAG output enable control for SysAD, SysADC, SysCmd and SysCmdP buses)

APPENDIX D CONSTRAINTS

The V_R4400MC has several constraints. For details, refer to V_R4000, V_R4400 USER'S MANUAL—ARCHITECTURE.

D