

IDT79R36100 Integrated RISController



**Hardware
User's Manual**



INTEGRATED DEVICE TECHNOLOGY, INC.



Integrated Device Technology, Inc.

IDT79R36100TM
Integrated RISControllerTM
Hardware User's Manual

Version 1.1
December 1995

2975 Stender Way, Santa Clara, California 95054
Telephone: (800) 345-7015 • TWX: 910-338-2070 • FAX: (408) 492-8674
Printed in U.S.A.
©1995 Integrated Device Technology, Inc.

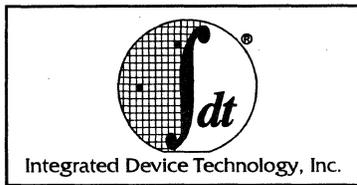
Integrated Device Technology, Inc. reserves the right to make changes to its products or specifications at any time, without notice, in order to improve design or performance and to supply the best possible product. IDT does not assume any responsibility for use of any circuitry described other than the circuitry embodied in an IDT product. The Company makes no representations that circuitry described herein is free from patent infringement or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, patent rights or other rights, of Integrated Device Technology, Inc.

LIFE SUPPORT POLICY

Integrated Device Technology's products are not authorized for use as critical components in life support devices or systems unless a specific written agreement pertaining to such intended use is executed between the manufacturer and an officer of IDT.

1. Life support devices or systems are devices or systems which (a) are intended for surgical implant into the body or (b) support or sustain life and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any components of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

The IDT logo is a registered trademark, and BiCameral, BurstRAM, BUSMUX, CacheRAM, DECnet, Double-Density, FASTX, Four-Port, FLEXI-CACHE, Flexi-PAK, Flow-thruEDC, IDT/c, IDTenvY, IDT/sae, IDT/sim, IDT/ux, MacStation, MICROSLICE, PalatteDAC, REAL8, R3041, R3051, R3052, R3071, R3081, R36100, R3721, R4600, R4650, R4700, RISCcontroller, RISCcore, RISC Subsystem, RISC Windows, SARAM, SmartLogic, SyncFIFO, SyncBiFIFO, SPC, TargetSystem and WideBus are trademarks of Integrated Device Technology, Inc. MIPS is a registered trademark, and RISCompiler, RISComponent, RISComputer, RISCware, RISC/os, R3000, and R3010 are trademarks of MIPS Computer Systems, Inc. Postscript is a registered trademark of Adobe Systems, Inc. AppleTalk, LocalTalk, and Macintosh are registered trademarks of Apple Computer, Inc. Centronics is a registered trademark of Genicom, Inc. Ethernet is a registered trademark of Digital Equipment Corp. PS2 is a registered trademark of IBM Corp.



About This Manual

This manual provides a description of the functional operation of the IDT79R36100 Integrated RISController™.

Summary of Contents

Chapter 1, "Overview," contains an overview of the R36100 microprocessor.

Chapter 2, "Instruction Set Architecture," contains an overview of the MIPS-1 architecture set and discusses the programmers' model for this device.

Chapter 3, "Cache Architecture," describes the fundamentals of general cache operations, as well as the particular organization of the on-chip caches of the R36100.

Chapter 4, "Virtual to Physical Address Translation and Address Map," describes the operating states of the processor, as well as the virtual to physical address translation mechanisms provided in the R36100.

Chapter 5, "Coprocessor 0 Register Set," describes the implementation of CPO found on the R36100, which are similar to those of the rest of the R30xx family.

Chapter 6, "Interrupt and Exception Handling," discusses exception handling issues in R36100-based systems, including software examples of exception handlers.

Chapter 7, "System Bus Interface Unit Overview," provides an overview of the operation of the execution core, as well as operation of the various memory controllers during both external transactions and internal peripheral transactions.

Chapter 8, "Memory Controller," provides an overview of the memory controller interface and a complete description of the signal pins and their timing.

Chapter 9, "I/O Controller," provides an overview of the I/O controller interface, a description of the signal pins and their timing, and a discussion of the relationship between the interface and typical hardware I/O devices.

Chapter 10, "DRAM Controller," provides an overview of the DRAM controller interface, a description of the signal pins and their timing, and a discussion of the relationship between the interface and typical external hardware DRAM systems.

Chapter 11, "Direct Memory Access (DMA) Controller," provides an overview of the DMA controller interface, a description of the signal pins and their timing, and a discussion of the relationship between the interface and typical internal and external hardware DMA systems.

Chapter 12, "Parallel Input/Output (PIO)," provides an overview of the PIO controller interface, a description of the signal pins and their timing, and a discussion of how PIOs relate to typical internal and external systems.

Chapter 13, "Peripheral Expansion Interrupt Controller," provides an overview of the PIO controller interface, a description of the signal pins and their timing, and a discussion of how expansion interrupts relate to typical internal and external systems.

Chapter 14, "Timers," provides an overview of the Timer programming interface, a description of the signal pins and their waveforms, and a discussion of how the timers relate to typical internal and external systems.

Chapter 15, "Serial Ports," provides an overview of the serial port register interface, a description of the signal pins, and a discussion of various aspects of the signal timing.

Chapter 16, "Bidirectional Parallel Port," provides an overview of the bidirectional Centronics parallel port register interface, a description of the signal pins, and a discussion of various aspects of the signal timing.

Chapter 17, "Laser Printer Video Port," provides an overview of the laser printer video port register interface, a description of the signal pins, and a discussion of various aspects of the signal timing.

Chapter 18, "Reset Initialization and Input Clocking," discusses the reset initialization sequence required by the R36100, the configuration mode selectable features of the processor, and boot software requirements.

Chapter 19, "Debug Mode Features," discusses features that facilitate debugging of R36100-based systems.

For More Product Information

Details about the R36100 electrical interface can be found in the product's data sheet. Data sheets also include packaging and pin-out information.

For information about development tools, complementary support chips, and how to use this product in various applications, refer to IDT's online library of data sheets, applications notes, software reference manuals, and the IDT Advantage Program Guides.

Your local IDT sales representative can help you identify and use these resources.



Table of Contents

R36100 Device Overview	Chapter 1
Introduction	1 - 1
R36100 Feature List	1 - 3
Device Overview	1 - 4
CPU Core	1 - 4
System Control Co-Processor	1 - 4
Clock Generator Unit	1 - 4
Instruction Cache	1 - 5
Data Cache	1 - 5
Bus Interface Unit	1 - 5
Memory Controller	1 - 6
DRAM Controller	1 - 6
I/O Controller	1 - 6
DMA Control and Interface	1 - 6
Counter/Timers	1 - 7
PIO Interface	1 - 7
Serial Communications Controller	1 - 7
Interrupt Controller	1 - 7
IEEE 1284 Bi-directional Centronics	1 - 8
Laser Printer Video Interface	1 - 8
Logic Symbol	1 - 9
Pin Description	1 - 10
System Usage	1 - 14
Development Support	1 - 15
Performance Overview	1 - 16
Instruction Set Architecture	Chapter 2
Introduction	2 - 1
Processor Features Overview	2 - 1
CPU Registers Overview	2 - 2
Instruction Set Overview	2 - 2
Programming Model	2 - 5
Data Formats and Addressing	2 - 5
CPU General Registers	2 - 7
CPO Special Registers	2 - 7
Operating Modes	2 - 8
Pipeline Architecture	2 - 9
Pipeline Hazards	2 - 10
Instruction Set Summary	2 - 12
Instruction Formats	2 - 12
Instruction Notational Conventions	2 - 12
Load and Store Instructions	2 - 13
Big-Endian (32-bit memory system)	2 - 13
Little-Endian (32-bit memory system)	2 - 14
Big-Endian (16-bit memory system)	2 - 14
Little-Endian (16-bit memory system)	2 - 15
Computational Instructions	2 - 17
Jump and Branch instructions	2 - 20
Special Instructions	2 - 21
Co-processor Instructions	2 - 22
System Control Co-processor (CPO) Instructions	2 - 23
R36100 Opcode Encoding	2 - 24

Cache Architecture.....	Chapter 3
Introduction	3 - 1
Fundamentals of Cache Operation	3 - 1
R36100 Cache Organization	3 - 2
Basic Cache Operation.....	3 - 2
Memory Address to Cache Location Mapping.....	3 - 2
Cache Addressing	3 - 3
Write Policy	3 - 3
Partial Word Writes	3 - 3
Instruction Cache Line Size.....	3 - 4
Data Cache Line Size	3 - 4
Summary.....	3 - 5
Cache Operation.....	3 - 5
Basic Cache Fetch Operation	3 - 5
Cache Miss Processing.....	3 - 6
Instruction Streaming.....	3 - 6
Cacheable References	3 - 7
Software Directed Cache Operations	3 - 7
Cache Sizing	3 - 7
Cache Flushing.....	3 - 8
Forcing Data into the Caches	3 - 9
Cache-Locking Operation.....	3 - 10
Summary.....	3 - 13
Virtual-to-Physical Address Translation and Address Map	Chapter 4
Virtual Memory in the R3000A Architecture.....	4 - 1
Privilege States	4 - 2
User Mode Virtual Addressing.....	4 - 2
Kernel Mode Virtual Addressing.....	4 - 3
R36100 address translation.....	4 - 4
On-Chip Registers.....	4 - 6
Cache Miss Area	4 - 7
Summary.....	4 - 7
Coprocessor 0 Register Set	Chapter 5
Introduction	5 - 1
Coprocessor 0 Bus Interface Control.....	5 - 1
Cache Configuration Register.....	5 - 2
Lock ('Lock')	5 - 3
Reserved-High ('1')	5 - 3
Reserved-Low ('0').....	5 - 3
DBlockRefill ('DBR').....	5 - 3
D-CacheIndexControl ('DCI').....	5 - 3
Halt Mode ('Halt')	5 - 4
I-CacheIndexControl ('ICI')	5 - 4
ReduceFrequency ('RF')	5 - 4
ForceDCacheMiss ('FDCM')	5 - 5
ForceICacheMiss ('FICM')	5 - 5
DCacheWriteDisable('DWrD').....	5 - 6
I-CacheWriteDisable ('IWrd').....	5 - 6
The Cause Register	5 - 7
The EPC (Exception Program Counter) Register	5 - 8
Bad VAddr Register	5 - 8
The Status Register.....	5 - 8
PRId Register	5 - 11

Interrupt and Exception Handling	Chapter 6
Introduction	6 - 1
R36100 Exception Model	6 - 1
Precise vs. Imprecise Exceptions	6 - 2
Exception Processing	6 - 3
Exception Handling Registers	6 - 3
The Cause Register	6 - 3
The EPC (Exception Program Counter) Register	6 - 5
Bad VAddr Register	6 - 5
The Status Register	6 - 5
Exception Vector Locations	6 - 7
Exception Prioritization.....	6 - 7
Exception Latency	6 - 9
Interrupts Inputs in the R36100	6 - 10
Interrupt Operation in the R36100.....	6 - 10
Using the BrCond Inputs	6 - 12
Interrupt Handling.....	6 - 13
Interrupt Servicing.....	6 - 13
Basic Software Techniques For Handling Interrupts	6 - 14
Preserving Context.....	6 - 15
Determining The Cause Of The Exception	6 - 16
Returning From Exceptions	6 - 18
Special Techniques For Interrupt Handling	6 - 19
Interrupt Masking.....	6 - 19
Using BrCond For Fast Response	6 - 19
Cache Locking	6 - 22
Nested Interrupts.....	6 - 22
Catastrophic Exceptions	6 - 22
Handling Specific Exceptions	6 - 23
Address Error Exception	6 - 23
Breakpoint Exception	6 - 23
Bus Error Exception	6 - 24
Co-processor Unusable Exception.....	6 - 25
Interrupt Exception	6 - 25
Overflow Exception	6 - 26
Reserved Instruction Exception.....	6 - 26
Reset Exception	6 - 27
System Call Exception	6 - 28
System Bus Interface Unit Overview.....	Chapter 7
Introduction	7 - 1
Bus Interface Overview	7 - 1
Pin Description	7 - 2
System Bus Interface Signals	7 - 2
Clock and Reset Signals	7 - 3
Bus Interface Control Signals.....	7 - 3
CPU Core Transaction Types.....	7 - 5
Read Operation	7 - 5
Write Operations	7 - 6
Multiple Operations	7 - 6
Execution Engine Fundamentals	7 - 7
Execution Core Cycles.....	7 - 7
Cycles	7 - 7
Run Cycles.....	7 - 8
Stall Cycles	7 - 8

Internal Acknowledgment	7 - 9
Read Interface Timing Overview	7 - 9
Initiation of a Read Request	7 - 9
Memory Addressing.....	7 - 10
Initiation of the Data Phase.....	7 - 11
Bringing Data into the Processor.....	7 - 11
Terminating the Read.....	7 - 13
Latency Between Processor Operations	7 - 14
Processor Internal Activity.....	7 - 15
The Write Interface	7 - 17
Importance of Writes in R36100 Systems	7 - 17
Types of Write Transactions	7 - 17
32-Bit Write Transactions	7 - 18
16-Bit Transactions	7 - 18
8-Bit Transactions	7 - 19
Write Interface Timing Overview.....	7 - 19
Initiating the Write	7 - 19
Memory Addressing.....	7 - 20
The Data Phase.....	7 - 20
Terminating the Write	7 - 20
Write Buffer Full Operation	7 - 22
Memory Controller	Chapter 8
Introduction	8 - 1
Features	8 - 1
Block Diagram.....	8 - 1
Memory Bus Controller Interface Signals	8 - 3
Memory Controller Signals	8 - 3
MemCS(7:0)	8 - 3
IoCS(7:0).....	8 - 3
MemRdEnEven	8 - 3
MemRdEnOdd	8 - 4
MemWrEn(3:0)	8 - 4
MemByteEn(3:0)	8 - 4
MemAddr(29:26)	8 - 4
BIU Controller Signals	8 - 4
SysAddr(25:0)	8 - 5
SysData(31:0)	8 - 5
SysWait.....	8 - 5
Overview of the Memory Controller (MemCntrl)	8 - 5
Chip Selects	8 - 6
Transceiver Control Interface	8 - 6
Wait State Generator	8 - 7
Register Option Programmability	8 - 7
Register Descriptions	8 - 7
Memory MSB Base Address Register for Bank 7..0	8 - 9
Memory LSB Base Address Register for Bank 7..0	8 - 10
Memory MSB Bank Mask Register for Bank 7..0	8 - 10
Memory LSB Bank Mask Register for Bank 7..0	8 - 12
Memory and I/O Control Register for Bank 7..0	8 - 12
Port Size Width ('MemSize') Field	8 - 13
Memory LSB Wait State Register for Bank 7..0	8 - 14
Read Start Cycle to the First Datum Field	8 - 15
Write Start Cycle to the First Datum Field	8 - 15
Read Datum to DatumField.....	8 - 15
Write Datum to Datum Field	8 - 16
Memory MSB Wait State Register for Bank 7..0	8 - 16
Repeat Start Bus Cycle State 0 Field	8 - 16
Start of Read to AckN on Burst Reads Field.....	8 - 17

Byte Enables on Reads Field	8 - 17
Bus Turn-Around Field	8 - 17
Memory Controller Timing Diagrams	8 - 18
Read Transactions	8 - 18
Basic 1-Datum Read with 0 Wait-States	8 - 18
1-Datum Read with 0 Wait-States Using Odd Chip Select	8 - 20
Read with Wait-State Using Start Repeat Field	8 - 21
Read with Wait-State Using RdStart2Datum Field	8 - 22
Read with Wait-State Using SysWait	8 - 23
4-Word Burst Read with 0 Wait-States	8 - 24
Basic 16-bit PCMCIA-style I/O Read with 0 Wait-States	8 - 25
Write Transactions	8 - 27
Single Datum Write	8 - 27
1-Datum Write with 0 Wait-States Using FCT245-Type Field	8 - 29
1-Datum Write with Wait-State Using StartRepeat Field	8 - 30
1-Datum Write with Wait-State Using WrStart2Datum Field	8 - 31
1-Datum Write with Wait-State Using SysWait	8 - 32
Multi-Datum Burst Write	8 - 33
Multi-Datum Burst Write Using Wait-State with WrDatum2Datum	8 - 34
Basic PCMCIA-Type Memory Write with 0 Wait-States	8 - 35
Interleaved-Type Transactions	8 - 36
Interleaved Read Using FCT260-Type Field	8 - 36
Interleaved Read Using FCT245-Type Field	8 - 40
Interleaved Read Using FCT543-Type Field	8 - 41
Interleaved Writes	8 - 44
System Examples	8 - 46
32-bit SRAM Using 245 Transceivers	8 - 47
16-bit SRAM/ROM	8 - 51
8-bit SRAM/ROM	8 - 52
Dual-Port-Type	8 - 53
PCMCIA-Style Application	8 - 53
PCI-Style	8 - 54
I/O Controller	Chapter 9
Introduction	9 - 1
Features	9 - 1
Block Diagram	9 - 2
I/O Bus Controller Interface Signals	9 - 3
IoCS(7:0)	9 - 3
MemCS(7:0)	9 - 3
IoRd	9 - 3
IoDStrobe	9 - 3
IoWr	9 - 3
IoRdHWr	9 - 3
MemAddr(29:26)	9 - 4
MemWrEn(3:0)	9 - 4
MemByteEn(3:0)	9 - 4
BIU Controller Signals	9 - 4
SysAddr	9 - 4
SysData	9 - 4
SysWait	9 - 5
Overview of the I/O Controller	9 - 6
Chip Selects	9 - 6
Signal Control Interface	9 - 6
Wait State Generator	9 - 6
Register Option Programmability	9 - 7

Register Descriptions	9 - 8
Memory and I/O Control Register7..0	9 - 9
Memory Type ('MemType') Field	9 - 9
Portsize Width ('MemSize') Field	9 - 9
I-Type I/O Type	9 - 10
M-Type I/O Type	9 - 10
PCMCIA-I/O Style	9 - 10
I/O Controller Timing Diagrams	9 - 11
I/O Datum Size	9 - 11
Read Transactions	9 - 11
Basic I-Type I/O Read with 0 Wait-States	9 - 11
Basic I-Type I/O Write with 0 Wait-States	9 - 15
Basic M-Type I/O Write with 0 Wait-States	9 - 16
Read with Wait-State Using Start Repeat Field	9 - 18
Read with Wait-State Using RdStart2Datum Field	9 - 18
Read with Wait-State Using SysWait	9 - 19
1-Datum Write with Wait-State Using StartRepeat Field	9 - 20
1-Datum Write with Wait-State Using WrStart2Datum	9 - 20
1-Datum Write with Wait-State Using SysWait	9 - 21
System Examples	9 - 22
32-bit I/O Device Directly Connected to Bus	9 - 22
I/O Reset Application	9 - 22
32-bit I/O Device Using 245 Transceivers	9 - 23
32-bit I/O Device Using 543 Transceivers	9 - 24
Using More Than One Device Behind Each Transceiver	9 - 24
16-bit I/O Devices	9 - 25
8-bit I/O Devices	9 - 26
DRAM Controller	Chapter 10
Introduction	10 - 1
Features	10 - 1
Block Diagram	10 - 1
DRAM Bus Controller Interface Signals	10 - 3
DRAM Interface Signals	10 - 3
SysAddr(13:2)	10 - 3
DramRAS(3:0)	10 - 3
DramCAS(3:0)	10 - 3
DramWrEnEven	10 - 3
DramWrEnOdd	10 - 3
DramRdEnEven	10 - 4
DramRdEnOdd	10 - 4
BIU Controller Signals	10 - 4
SysDataOutput	10 - 4
SysWait	10 - 4
Overview of the DRAM Controller	10 - 4
Address Mapping	10 - 5
32-bit and 16-bit Mode Support	10 - 5
Types of Memory Supported	10 - 6
Programmable Wait State Generation	10 - 6
Page Comparator Algorithm	10 - 6
Unaligned Page Accesses	10 - 7
Refresh Timing	10 - 7
Initialization	10 - 7
Programmable Features	10 - 7
Signal Control Interface	10 - 7
Wait State Generator	10 - 7
Register Option Field Programmability	10 - 8
Register Descriptions	10 - 8
DRAM Refresh Count Register	10 - 9
Staggered Refresh	10 - 9

Refresh Arbitration	10 - 9
Panic Mode Refresh Application	10 - 9
Reduced Frequency Mode Application	10 - 10
DRAM Refresh Compare Register	10 - 10
DRAM RAS Multiplexer Select Register for Pair.....	10 - 11
DRAM CAS Multiplexer Select Register for Pair.....	10 - 11
DRAM MSB Bank Mask Register for Bank 0..3.....	10 - 15
DRAM LSB Control Register for Bank 0..3.....	10 - 15
RAS Bank Mask ('RASPageMask') Field.....	10 - 16
DRAM Type ('DramType') Field	10 - 16
FCT543-Type (Latched Non-Multiplexer Type)	10 - 17
FCT245 Type (Non-latched Transceiver Type)	10 - 17
FCT260-Type (Latched Multiplexer Type).....	10 - 17
Port Size ('Size') Field	10 - 18
Page Type ('PageType') Field	10 - 18
DRAM MSB Control Register for Bank 0..3	10 - 18
RAS Precharge Period ('RASPre') Field	10 - 19
RAS Address Hold Time ('RASAddrHold') Field.....	10 - 19
Address Setup Time to RAS and to CAS Field	10 - 19
CAS Active Pulse Width Field	10 - 20
DRAM Read Cycle Bus Turn-Around Field	10 - 20
DRAM Write Cycle Bus Turn Around Field	10 - 20
Burst Acknowledge Placement Field	10 - 21
Timing Diagrams	10 - 21
Standard DRAM Chip Summary.....	10 - 22
Basic New Page DRAM Read.....	10 - 22
RAS Asserted at End of Transfer	10 - 23
RAS Asserted at Start of Transfer	10 - 24
RAS Asserted Throughout Transfer	10 - 24
RAS Precharge Field.....	10 - 26
RAS Address Hold Field	10 - 27
Address Setup Field.....	10 - 28
CAS Width Field.....	10 - 29
Multiple Datum Reads	10 - 30
Basic DRAM Write	10 - 31
RAS Asserted at Start of Write.....	10 - 32
RAS Asserted at End of Write	10 - 33
RAS Asserted Throughout Write.....	10 - 34
Other DRAM timing Controls	10 - 35
Write Bus Turn-Around	10 - 35
Two Datum Write Transaction.....	10 - 36
Interleaved Reads	10 - 37
Interleaved FCT245 Reads.....	10 - 37
Interleaved FCT260 Reads.....	10 - 38
Interleaved FCT543 Reads.....	10 - 41
Interleaved Writes	10 - 44
Single Word Interleaved FCT245 Write.....	10 - 44
Interleaved FCT245 Writes	10 - 44
Single Word Interleaved FCT260 Write.....	10 - 46
Interleaved FCT260 Writes	10 - 46
Interleaved FCT543 Writes	10 - 46
Refresh	10 - 48
System Examples	10 - 49
DRAM System Using FCT245 Transceivers.....	10 - 49
Low Cost DRAM System Using FCT245 Transceivers.....	10 - 50
Very Low Cost DRAM System without Transceivers	10 - 50
DRAM System Using FCT260 Multiplexers	10 - 50
DRAM System Using FCT543 Registered Transceivers.....	10 - 51

Direct Memory Access (DMA) Controller	Chapter 11
Introduction	11 - 1
Features	11 - 1
Block Diagram	11 - 2
Overview	11 - 2
Internal DMA Channels	11 - 2
Internal DMA Algorithm	11 - 5
External DMA Channels	11 - 6
Pin Descriptions	11 - 7
Direct Memory Access (DMA) Controller Signals	11 - 7
DmaBusReq(1:0)	11 - 7
DmaBusGnt(1:0)	11 - 7
DmaDone	11 - 7
System Control Signals Used During DMA Controller	
Accesses	11 - 8
SysALEn	11 - 8
SysBurstFrame	11 - 8
SysRd	11 - 8
SysWr	11 - 8
MemWrEn(3:0)	11 - 8
SysRd	11 - 8
MemWrEn()	11 - 8
SysDataRdy	11 - 8
Register Descriptions	11 - 9
Internal DMA Controller Register Descriptions	11 - 9
DMA LSB Source Address Register for Channel 0..3	11 - 10
DMA LSB Source Address Register for Link A..D	11 - 10
DMA MSB Source Address Register for Channel 0..3	11 - 10
DMA MSB Source Address Register for Link A..D	11 - 10
DMA LSB Target Address Register for Channel 0..3	11 - 10
DMA LSB Target Address Register for Link A..D	11 - 10
DMA MSB Target Address Register for Channel 0..3	11 - 11
DMA MSB Target Address Register for Link A..D	11 - 11
DMA LSB Count Register for Channel 0..3	11 - 11
DMA LSB Count Register for Link A..D	11 - 11
DMA MSB Count Register for Channel 0..3	11 - 11
DMA MSB Count Register for Link A..D	11 - 11
DMA LSB Control Register for Channel 0..3	11 - 12
DMA LSB Control Register for Link A..D	11 - 12
Arbitration Type ('Arb') Field	11 - 12
Allow DMADone ('Done') Field	11 - 13
Wait for Interrupt ('WInt') Field	11 - 13
Burst Type ('Burst') Field	11 - 13
Cacheable ('Cache') Field	11 - 14
Source Byte Enable Type ('SBE') Field	11 - 14
Target Byte Enable Type ('TBE') Field	11 - 14
Source Endianness Type ('SEndian') Field	11 - 14
Target Endianness Type ('TEndian') Field	11 - 15
Increment Source Address ('SInc') Field	11 - 15
Increment Target Address ('TInc') Field	11 - 15
DMA MSB Control Register for Channel 0..3	11 - 15
DMA MSB Control Register for Link A..D	11 - 15
Stop ('Stop') Field	11 - 16
Break ('Break') Field	11 - 16
ReservedLink Field ('RsvdLink')	11 - 16
Link ('Link') Field	11 - 16
External DMA Controller Register Descriptions	11 - 17
External DMA Control Register 0..1	11 - 17
Stop Channel ('EC') Field	11 - 18
Bus Request Protocol High ('ReqH') Field	11 - 18

Bus Grant Protocol High ('GntH') Field	11 - 18
Sample MemWrEn and SysBurstFrame 1 Clock Later ('SampleLate') Field.....	11 - 18
System Examples	11 - 22
Memory-to-memory Copying	11 - 22
Transfers between I/O and Memory	11 - 22
Distinguishing Between CPU and Internal DMA Accesses.....	11 - 23
Parallel Input/Output (PIO).....	Chapter 12
Introduction	12 - 1
Features	12 - 1
Block Diagram.....	12 - 1
Overview.....	12 - 1
Pin Descriptions	12 - 1
PIO(n)	12 - 2
Register Definitions.....	12 - 3
PIO Data Register 0..2.....	12 - 3
PIO Data ('PIOData') Field	12 - 3
PIO Direction Register 0..2	12 - 3
Lock ('Lock') Field.....	12 - 4
Direction ('Dir') Field	12 - 4
PIO Effect Select Register 0..2	12 - 4
Lock ('Lock') Field.....	12 - 4
Effect Select ('EffectSel') Field.....	12 - 5
Peripheral Expansion Interrupt Controller	Chapter 13
Introduction	13 - 1
Features	13 - 1
Block Diagrams	13 - 1
Overview.....	13 - 2
Pin Descriptions	13 - 2
Exception Signals.....	13 - 2
ExcSInt(2:0)	13 - 2
ExcInt(4:3)	13 - 2
ExcSBrCond(3:2).....	13 - 2
Expansion Interrupt Mask Register 0..1	13 - 3
Expansion Interrupt Pending Register 0..1	13 - 3
Reserved Low ('0') Field	13 - 3
Mask Bits ('Mask') Field.....	13 - 4
Pending Bits ('Pend') Field	13 - 4
Expansion Interrupt DMA Select Register	13 - 5
Select Interrupt 'SelInt()' Field	13 - 5
Timers.....	Chapter 14
Introduction	14 - 1
Features	14 - 1
Block Diagram.....	14 - 1
Overview.....	14 - 2
Pin Descriptions	14 - 2
Timer Peripheral Signals	14 - 2
TC(2:0).....	14 - 2
TimerGate(2:0)	14 - 2
Register Descriptions	14 - 3
Timer Prescaler Count Register	14 - 3
Timer Count Register 0..2	14 - 3
Timer Compare Register 0..2.....	14 - 4
Timer Pulse Width Modulation Register 0.....	14 - 4
Timer Control Register 0..2	14 - 4
Lock ('Lock') Field.....	14 - 5

Lock Count and Compare ('LockCC') Field	14 - 5
Write Compare Ack ('Ack') Field	14 - 5
PIO is Output TC ('TC') Field	14 - 6
PIO is Input Gate ('Gate') Field	14 - 6
BusTimeout ('BTO') Field.....	14 - 6
Timer Enable ('TimerEn') Field	14 - 6
Prescaler Select ('PSel') Field	14 - 6
Serial Ports.....	Chapter 15
Introduction	15 - 1
Features	15 - 1
Block Diagram.....	15 - 1
Overview.....	15 - 2
DMA Connections	15 - 2
DMA Data Movement	15 - 3
Serial Data Speed Calculations	15 - 3
Compatibility Issues.....	15 - 4
Pin Definitions	15 - 5
Serial Port Signals.....	15 - 5
Register Definitions.....	15 - 7
Data Register Usage	15 - 8
MetaRegister Usage.....	15 - 8
MetaMetaRegister Usage	15 - 9
Serial LSB Status Register	15 - 9
LSB Status Bits	15 - 9
Serial MSB Status Register	15 - 10
Serial Interrupt Vector.....	15 - 10
Serial LSB Sync Character/Address Register	15 - 12
Serial MSB Sync Character/Address Register	15 - 12
Receiver Data FIFO Register.....	15 - 12
Serial Clock Status Register	15 - 12
Serial LSB Baud Rate Generator Compare Register	15 - 13
Serial MSB Interrupt Enable Register	15 - 13
Meta Pointer Register	15 - 14
Serial Meta Register	15 - 14
Serial LSB Interrupt Enable Register	15 - 15
Receiver Interrupt Enable Mode	15 - 16
Parity Interrupt Enable Mode	15 - 16
Transmitter Interrupt Enable Mode.....	15 - 16
External Status Interrupt Enable	15 - 16
Serial Interrupt Vector Register	15 - 17
Serial Receiver Control Register.....	15 - 17
Receiver Bits per Character	15 - 18
Auto-Handshaking	15 - 18
Sync Char Hunt Mode.....	15 - 18
CRC Checking	15 - 18
Address Match Search	15 - 19
Address Multiple Match Mask	15 - 19
Receiver Enable	15 - 19
Serial LSB Mode Register	15 - 19
Clock Sampling Mode.....	15 - 20
Sync Character Mode	15 - 20
Stop Bit Mode	15 - 20
Parity Mode.....	15 - 20
Parity Enable	15 - 21
Serial Transmitter Control Register	15 - 21
DTR Status	15 - 21
Transmitter Bits per Character	15 - 22
Transmitter Break.....	15 - 22
Transmitter Enable	15 - 22
CRC Type.....	15 - 22

RTS Status	15 - 22
Transmitter CRC Enable	15 - 22
Serial LSB Sync Character/Address Register.....	15 - 23
Serial MSB Sync Character/Address Register.....	15 - 23
Transmitter Data FIFO Register	15 - 24
Serial Reset Register	15 - 24
Interrupt Enable	15 - 24
Serial MSB Mode Register	15 - 24
CRC Preload	15 - 25
Endec Mode	15 - 25
Automatic SDLC Poll Frame Mode	15 - 25
Automatic Transmitter SDLC Idle Mark.....	15 - 25
Automatic Transmitter Underrun Abort.....	15 - 26
SDLC Loop Mode.....	15 - 26
Sync Char Length	15 - 26
Serial LSB Clock Control Register	15 - 26
Receiver Clock Input Source	15 - 27
Transmitter Clock Input Source.....	15 - 27
SerialSecondaryClk() Direction	15 - 27
SerialSecondaryClk() Output Clock Select	15 - 27
Serial LSB Baud Rate Clock Gen Compare Register.....	15 - 28
Serial MSB Baud Rate Clock Gen Compare Register.....	15 - 28
PLL Control Register	15 - 29
PLL Mode	15 - 29
Loopback Test Enable	15 - 29
Echo Test Enable	15 - 30
Baud Rate Clock Generator Clock Input.....	15 - 30
Baud Rate Clock Generator Enable	15 - 30
Serial MSB Interrupt Enable Register	15 - 30
Serial MetaMetaRegister.....	15 - 30
Interrupt Enable Bits	15 - 31
MetaMeta Control Register Enable	15 - 31
Serial LSB Frame Status FIFO Data Register	15 - 32
Serial MSB Frame Status FIFO Data Register.....	15 - 32
Frame Status FIFO Overflow Flag.....	15 - 32
Frame Status Not Empty Flag	15 - 32
Count	15 - 32
Serial Meta Meta Register.....	15 - 33
Wait for Receiver CRC	15 - 33
Automatically Drive SDLC SerialTxData High When Disabled.....	15 - 33
Automatically De-assert SerialRTS After SDLC Flag	15 - 34
Automatically Reset Transmitter Underrun	15 - 35
Automatic Insertion of Transmitter Opening Flag	15 - 35
Bidirectional Parallel Port.....	Chapter 16
Introduction	16 - 1
Features	16 - 1
Block Diagram.....	16 - 1
Overview.....	16 - 2
Negotiation Phase	16 - 4
Nibble Mode Phase.....	16 - 4
Byte Mode Phase.....	16 - 4
Extended Capabilities Port (ECP) Mode Phase.....	16 - 4
Enhanced Parallel Port (EPP) Mode Phases.....	16 - 5
CPU Control Reserved Mode Phases	16 - 5
Programmable Timing.....	16 - 5
Pin Descriptions	16 - 6
Bidirectional Parallel Port Centronics Interface Signals	16 - 6
CentStrobe.....	16 - 6
CentAck.....	16 - 6

CentBusy	16 - 6
CentPaperError	16 - 6
CentSelect	16 - 7
Centronics Select	16 - 7
CentAutoFeed	16 - 7
Centronics Auto Feed	16 - 7
CentInit	16 - 7
Centronics Initialize	16 - 7
CentFault	16 - 7
Centronics Fault	16 - 7
CentSelectIn	16 - 7
Centronics Select Input	16 - 7
Bidirectional Parallel Port Centronics Peripheral and Host Interface	
Signals	16 - 8
CentCS	16 - 8
CentWrStrobe	16 - 8
CentRdOEn	16 - 8
CentHostStrobe	16 - 8
CentHostOEn	16 - 8
Register Definitions	16 - 8
Centronics Sub Mode Register	16 - 9
Centronics Compatible Sub Modes Field	16 - 9
Centronics Status Register	16 - 9
IoCS(7), IoCS(6) Mask Enable	16 - 10
Centronics Negotiation Interrupt Pending	16 - 10
Centronics Negotiation Interrupt Enable	16 - 11
Centronics Reset Interrupt Pending	16 - 11
Centronics Reset Interrupt Enable	16 - 11
Printer Error Field	16 - 11
Printer On Line Select Field	16 - 11
Printer Fault Field	16 - 11
Printer Acknowledge Negated Field	16 - 12
Printer Busy Field	16 - 12
Centronics Control Register	16 - 12
Negotiation XFlag Reply Field	16 - 12
Negotiation Mode Field	16 - 13
Centronics Nibble Data Register	16 - 13
Centronics Host Status Register	16 - 13
AutoFeed Negated Field	16 - 14
Initialize Negated Field	16 - 14
Select In Negated Field	16 - 14
Host Strobe Negated Field	16 - 14
Centronics Minimum Delay Register	16 - 14
250ns Delay Type Field	16 - 15
500ns Delay Type Field	16 - 15
Timing Diagram	16 - 16
System Example	16 - 17
Laser Printer Video Port	Chapter 17
Introduction	17 - 1
Features	17 - 1
Block Diagram	17 - 1
Overview	17 - 2
FIFO Depth	17 - 2
Internal PLL Frequency	17 - 3
Reset Restriction	17 - 3
Interrupt Descriptions	17 - 3
LaserFIFONotFullInt	17 - 3
LaserPageInt	17 - 3
LaserBandInt	17 - 3
LaserLineInt	17 - 3

Pin Descriptions	17 - 3
Laser Printer Video Control Interface Signals	17 - 3
Laser Printer Video Interface Signals	17 - 3
LaserVideoClkIn	17 - 3
LaserVideoData	17 - 3
LaserLineSync	17 - 4
LaserPageSync	17 - 4
Register Definitions	17 - 5
Laser Video Control Register	17 - 5
Video FIFO Full	17 - 5
Request Size	17 - 6
Reset FIFO	17 - 6
Last Band Field	17 - 6
Video Direction Field	17 - 6
Inverse Video Data Field	17 - 6
Phase Lock Loop Enabled Field	17 - 7
Laser Video Vertical Skip Register	17 - 7
Laser Vertical Skip Dis Field	17 - 7
Vertical Skip Count Field	17 - 7
Laser Video Horizontal Skip Register	17 - 8
Horizontal Skip Enable Field	17 - 8
Horizontal Skip Count Field	17 - 8
Laser Video Band Line Count Register	17 - 9
Band Line (Lines per Band/Page) Count Field	17 - 9
Laser Video Horizontal Count Register	17 - 9
Laser Video Data Register	17 - 10
System Example	17 - 11
Reset Initialization and Input Clocking	Chapter 18
Introduction	18 - 1
Reset Timing	18 - 1
Reset Configuration Mode Features	18 - 1
Reset Configuration Mode Pin Descriptions	18 - 1
Exception Signals	18 - 1
SysReset	18 - 1
ExcSInt(2:0)	18 - 1
BigEndian	18 - 2
BootProm8	18 - 2
BootProm16	18 - 2
R3000A Equivalent Modes	18 - 2
Reset Behavior	18 - 2
Boot Software Requirements	18 - 3
Detailed Reset Timing Diagrams	18 - 3
Reset Pulse Width	18 - 4
Mode Initialization Timing Requirements	18 - 4
Reset Setup Time Requirements	18 - 6
ClkIn Requirements	18 - 6
Debug Mode Features	Chapter 19
Introduction	19 - 1
Features	19 - 1
Tri-State-able Outputs	19 - 1
Tracepoint Registers	19 - 1
Extended CP0 Cache Configuration Register	19 - 2
Cause and EPC Register Writes	19 - 2
JTAG Scan Path	19 - 2
Features Specific to Debug/Emulators	19 - 2

Pin Descriptions	19 - 2
Debug/Emulator and Diagnostic Signals.....	19 - 2
DiagCache/UnCache.....	19 - 2
DiagInst/Data.....	19 - 2
DiagRun	19 - 2
DiagBranchTaken	19 - 2
DiagJRorExe	19 - 3
DiagInternalWr	19 - 3
DiagTriState.....	19 - 3
DiagInstCacheWrDis	19 - 3
DiagFCM.....	19 - 3
DiagIntDis	19 - 3
DiagNoCS	19 - 3
DiagInternalDmaBusGnt.....	19 - 3
JTAG Signals	19 - 4
JtagClkIn (TCLK).....	19 - 4
JtagModeSelect (TMODE)	19 - 4
JtagDataIn (TDI)	19 - 4
JtagDataOut (TDO)	19 - 4
JtagReset (TRES*)	19 - 4
Register Descriptions	19 - 4
MSB Debug Tracepoint Address Register.....	19 - 4
LSB Debug Tracepoint Address Register.....	19 - 4
Debug Tracepoint Control Register	19 - 5
Reserved Low ('0') Field	19 - 5
Cause is Tracepoint ('CTP') Field	19 - 5
Tracepoint ('TP') Field.....	19 - 6
Debug Control Register	19 - 6
Reserved Low ('0') Field	19 - 6
Writability ('Wr') Field	19 - 6
Initializing SysClk for Test.....	19 - 7
Using Diag for Instruction Disassembly.....	19 - 7



Integrated Device Technology, Inc.

List of Tables

Table No.	Table Title	Page
Table 1.1	R36100 Pin Descriptions	1 - 10
Table 2.1	Instruction Set Mnemonics	2 - 4
Table 2.2	R36100 CPO Registers	2 - 8
Table 2.3	Big-Endian (32-bit memory system)	2 - 13
Table 2.4	Byte Addressing in Load/Store Operations (32-bit memory)	2 - 14
Table 2.5	Big-Endian (16-bit memory system)	2 - 14
Table 2.6	Byte Addressing in Load/Store Operations (16-bit memory)	2 - 15
Table 2.7	Load and Store Instructions	2 - 16
Table 2.8	ALU Immediate Operations	2 - 17
Table 2.9	Three Operand Register-Type Operations	2 - 18
Table 2.10	Shift Operations	2 - 19
Table 2.11	Multiply and Divide Operations	2 - 19
Table 2.12	Jump Instructions	2 - 20
Table 2.13	Branch Instructions	2 - 21
Table 2.14	Special Instructions	2 - 21
Table 2.15	Co-Processor Operations	2 - 22
Table 2.16	System Control Co-Processor (CPO) Operations	2 - 23
Table 2.17	Opcode Encoding	2 - 24
Table 3.1	Instruction Cache to Address Mapping under Various Cache Locking Conditions	3 - 13
Table 3.2	Data Cache to Address Mapping under Various Cache Locking Conditions	3 - 13
Table 4.1	Virtual and Physical Address Relationships in Base Versions	4 - 5
Table 4.2	R36100 On-Chip Resources and Address Map.....	4 - 6
Table 4.3	Example: FIFO load code using FCM memory space.	4 - 7
Table 5.1	R36100 CPO Register Addresses	5 - 2
Table 5.2	R36100 Cache Configuration Register Lock Field	5 - 3
Table 5.3	R36100 DBlockRefill Field	5 - 3
Table 5.4	R36100 D-Cache Index Control Field	5 - 3
Table 5.5	R36100 Halt Field	5 - 4
Table 5.6	R36100 I-Cache Index Control Field	5 - 4
Table 5.7	R36100 Reduced Frequency Mode Field	5 - 4
Table 5.8	R36100 ForceDCacheMiss Field	5 - 5
Table 5.9	R36100 ForceICacheMiss Field	5 - 5
Table 5.10	R36100 Data Cache Write Disable Field	5 - 6
Table 5.11	R36100 Instruction Cache Write Disable Field	5 - 6
Table 5.12	Cause Register Exception Codes	5 - 7
Table 6.1	R3000 Architecture Exceptions	6 - 2
Table 6.2	Cause Register Exception Codes	6 - 5
Table 6.3	Exception Vectors When BEV = 0	6 - 7
Table 6.4	Exception Vectors When BEV = 1	6 - 7
Table 6.5	R36100 Exception Priority	6 - 9
Table 8.1	List of the Memory and I/O Controller Registers (1 of 2).	8 - 8
Table 8.2	List of the Memory and I/O Controller Registers (2 of 2)	8 - 9
Table 8.3	Memory and I/O Controller Base Addresses.	8 - 10
Table 8.4	Memory Mask Field Encoding	8 - 11
Table 8.5	Memory and I/O Control Register Bit Assignments.	8 - 12
Table 8.6	Memory Type Field ('MemType') Encoding.....	8 - 13
Table 8.7	PortSize ('MemSize') Encoding	8 - 13
Table 8.8	Memory LSB Wait State Register ('MemLSBWaitStateReg') Bit Assignments	8 - 14
Table 8.9	Start to the first Datum ('RdStart2Datum' and 'WrStart2 Datum') Field Encoding.	8 - 15

Table 8.10	Datum-to-Datum (RdDatum2Datum, WrDatum2Datum) Field Encoding	8 - 15
Table 8.11	Memory MSB Wait State Register ('MemMSBWaitStateReg') Bit Assignments	8 - 16
Table 8.12	Repeat Start Bus Cycle State 0 ('StartRepeat') Field Encoding	8 - 16
Table 8.13	First Read to AckN on Burst Reads ('Start2BurstAck') Field Encoding	8 - 17
Table 8.14	Byte Enables on Reads ('BEn') Field Encoding	8 - 17
Table 8.15	Bus Turn-Around ('BTA') Field Encoding	8 - 17
Table 8.16	PCMCIA and R36100 Functional Equivalents	8 - 53
Table 8.17	PCI and R36100 Functional Equivalents	8 - 54
Table 9.1	Memory and I/O Controller Register Addresses and Descrip- tions	9 - 8
Table 9.2	MemCntrlReg Memory-Type Field ('MemType')	9 - 9
Table 9.3	PortSize ('MemSize') Encoding	9 - 9
Table 10.1	DRAM Controller Registers	10 - 8
Table 10.2	DRAM Refresh Count Register ('DramRefreshCountReg') Bit Assignments	10 - 9
Table 10.3	DRAM Refresh Compare Register ('DramRefreshCompare Reg') Bit Assignments	10 - 10
Table 10.4	Refresh Disable ('RefreshDis') Field Encodings	10 - 10
Table 10.5	Common Refresh Settings for 8ms/512 or 16ms/1024 DRAMs	10 - 11
Table 10.6	DRAM RAS Mux Select Register Bit Assignments	10 - 11
Table 10.7	DRAM CAS Mux Select Register ('DramCasMuxSelReg') Bit Assignments	10 - 12
Table 10.8	Example 'DramRasMuxSelReg' and 'DramCasMuxSelReg' Settings	10 - 13
Table 10.9	Example Bank Base Address Register ('DramMSBBase AddrReg') Assignment	10 - 14
Table 10.10	Example Bank Base Address Register ('DramMSBBase AddrReg') Assignment	10 - 14
Table 10.11	DRAM MSB Bank Mask Bit Settings	10 - 15
Table 10.12	DRAM LSB Control Register ('DramLSBControlReg') Bit Assignments	10 - 15
Table 10.13	Page Mask ('PMask') Bits	10 - 16
Table 10.14	DRAM LSB Page Mask Bit Settings	10 - 16
Table 10.15	DRAM Type ('DramType') Settings	10 - 16
Table 10.16	DRAM Port Width ('Size') Encoding Field	10 - 18
Table 10.17	Page Type ('PageType') (bits 2:0) Field	10 - 18
Table 10.18	DRAM MSB Control Register Bit Assignments	10 - 18
Table 10.19	RAS Precharge ('RASP') Field Encodings	10 - 19
Table 10.20	RAS Address Hold Time ('RASAddrHold') Field Encoding	10 - 19
Table 10.21	Address Setup Time to RAS or to CAS ('AddrSetup') Field Encoding	10 - 20
Table 10.22	CAS Width ('CASW') Field Encoding	10 - 20
Table 10.23	DRAM Read Cycle Bus Turn-Around ('DramRdBTA') Field Encoding	10 - 20
Table 10.24	DRAM Write Cycle Bus Turn-Around ('DramWrBTA') Field Encoding	10 - 21
Table 10.25	Burst Read Acknowledge ('DramBurstAck') Encoding	10 - 21
Table 11.1	Fixed Priority Encoding	11 - 3
Table 11.2	Internal Channel DMA Controller Register Address Map	11 - 9
Table 11.3	Internal DMA LSB Control Register ('DmaLSBControl Reg') Bit Assignments	11 - 12
Table 11.4	Arbitration Type ('Arb') Field Encoding	11 - 12
Table 11.5	Keep Bus ('Bus') Field Encoding	11 - 13
Table 11.6	Allow DMADone ('Done') Field Encoding	11 - 13
Table 11.7	Wait for Interrupt ('Int') Field Encoding	11 - 13

Table 11.8	Burst Type ('Burst') Field Encoding	11 - 13
Table 11.9	Cacheable ('Cache') Field Encoding	11 - 14
Table 11.10	Source Byte Enable Type ('SBE') Field Encoding	11 - 14
Table 11.11	Target Byte Enable Type ('TBE') Field Encoding	11 - 14
Table 11.12	Source Big Endianess Type ('SEndian') Field Encoding	11 - 14
Table 11.13	Target Big Endianess Type ('TEndian') Field Encoding	11 - 15
Table 11.14	Increment Source Address ('HInc') Field Encoding	11 - 15
Table 11.15	Increment Target Address ('TInc') Field Encoding	11 - 15
Table 11.16	Internal DMA MSB Control Register ('DmaMSBControl Reg') Bit Assignments	11 - 16
Table 11.17	Stop ('Stop') Field Encoding	11 - 16
Table 11.18	Break ('Break') Field Encoding	11 - 16
Table 11.19	Reserved Link ('RsvdLink') Field Encoding	11 - 16
Table 11.20	Link ('Link') Field Encoding	11 - 16
Table 11.21	External DMA Controller Register Address Assign- ments	11 - 17
Table 11.22	External DMA Control Register ('ExtDmaControlReg') Bit Assignments	11 - 17
Table 11.23	Enable Channel ('EC') Field Encoding	11 - 18
Table 11.24	Bus Request Protocol High ('ReqH') Field Encoding	11 - 18
Table 11.25	Bus Grant Protocol High ('GntH') Field Encoding	11 - 18
Table 11.26	Sample MemWrEn and SysBurstFrame 1 Clock Later ('SampleLate') Field Encoding	11 - 18
Table 12.1	Alternate R36100 functions mapped to PIO pins	12 - 2
Table 12.2	PIO Register Address Assignments	12 - 3
Table 12.3	PIO Data ('PIOData') Field Encoding	12 - 3
Table 12.4	Lock ('Lock') Field Encoding	12 - 4
Table 12.5	Direction ('Dir') Field Encoding	12 - 4
Table 12.6	Lock ('Lock') Field Encoding	12 - 4
Table 12.7	Effect Select ('EffectSel') Field Encoding	12 - 5
Table 13.1	Expansion Interrupt Controller Register Address Assign- ments	13 - 3
Table 13.2	Expansion Interrupt Mask Register 1 and Expansion Interrupt Pending Register 1 Bit Assignments	13 - 4
Table 13.3	Expansion Interrupt Mask Register 0 and Expansion Interrupt Pending Register 0 Bit Assignments	13 - 4
Table 13.4	Pending Interrupt Field Encoding	13 - 4
Table 13.5	Interrupt Mask Field Encoding	13 - 5
Table 13.6	DMA Channel versus Interrupt De-Multiplexer	13 - 5
Table 13.7	Select Interrupt ('SelInt()') Field Encoding	13 - 5
Table 14.1	Timer Register Physical Address Map	14 - 3
Table 14.2	Timer Pulse Width Modulation Register ('TimerPWMReg') Bit Fields	14 - 4
Table 14.3	Timer Control Register ('TimerControlReg') Bit Assign- ments	14 - 5
Table 14.4	Lock ('Lock') Field Encoding	14 - 5
Table 14.5	Lock Count and Compare ('LockCC') Field Encoding	14 - 5
Table 14.6	Write Compare Ack ('Ack') Field Encoding	14 - 5
Table 14.7	PIO is Output TC ('TC') Field Encoding	14 - 6
Table 14.8	PIO is Input Gate ('Gate') Field Encoding	14 - 6
Table 14.9	BusTimeout ('BTO') Field Encoding	14 - 6
Table 14.10	Timer Enable ('TimerEn') Field Encoding	14 - 6
Table 14.11	Prescaler Select ('PSel') Field Encoding	14 - 6
Table 15.1	List of Serial Port MetaRegisters	15 - 7
Table 15.2	Serial Port Internal Read Registers for each Channel (as pointed to by the MetaRegister)	15 - 7
Table 15.3	Serial Port Internal Write Registers for each Channel (as pointed to by the Meta Register)	15 - 8
Table 15.4	Serial Port Internal Read Registers for each Channel (as pointed to by the MetaPtrRegister 0xF)	15 - 8
Table 15.5	Serial Port Internal Write Register for each Channel (as pointed to by the MetaPtrRegister 0xOF)	15 - 8

Table 15.6	Serial LSB Status Register ('SerialLSBStatusReg') Bit Assignments	15 - 9
Table 15.7	LSB Status Bit Field Encoding	15 - 9
Table 15.8	Serial MSB Status Register ('SerialMSBStatusReg') Field Encoding	15 - 10
Table 15.9	MSB Status Bit Field Encoding	15 - 10
Table 15.10	Serial Interrupt Vector Register ('SerialRxControlReg') Bit Assignments	15 - 10
Table 15.11	Interrupt Vector ('IntVector') Field Encoding	15 - 11
Table 15.12	Serial Clock Status Register ('SerialClkStatusReg') Bit Assignments	15 - 13
Table 15.13	Clock Status Bit Field Encoding	15 - 13
Table 15.14	Serial MSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments	15 - 14
Table 15.15	Serial Meta Register ('SerialMetaRegister') Bit Assign- ments	15 - 14
Table 15.16	Reset A Command ('ResetA') Field Encoding	15 - 15
Table 15.17	Reset B Command ('ResetB') Field Encoding	15 - 15
Table 15.18	Pointer Value ('Pointer') Field Encoding	15 - 15
Table 15.19	Serial LSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments	15 - 16
Table 15.20	Receiver Interrupt Enable Mode ('RxIntEn') Field Encoding	15 - 16
Table 15.21	Parity Interrupt Enable ('ParityIntEn') Field Encoding	15 - 16
Table 15.22	Transmitter Interrupt Enable Mode ('TxIntEn') Field Encoding	15 - 16
Table 15.23	External Status Interrupt Enable ('ExtStatusIntEn') Field Encoding	15 - 16
Table 15.24	Serial Interrupt Vector Register ('SerialIntVectorReg') Bit Assignments	15 - 17
Table 15.25	Interrupt Vector Bit Encoding	15 - 17
Table 15.26	Serial Receiver Control Register ('SerialRxControlReg') Bit Assignments	15 - 18
Table 15.27	Receiver Bits per Character ('RxBitsPerChar') Field Encoding	15 - 18
Table 15.28	Auto-Handshaking ('AutoHandShake') Field Encoding	15 - 18
Table 15.29	Sync Char Hunt ('SyncHunt') Mode Field Encoding	15 - 18
Table 15.30	CRC Checking ('CRCCheck') Field Encoding	15 - 18
Table 15.31	Address Match Search ('AddrMatch') Field Encoding	15 - 19
Table 15.32	Address Multiple Match Mask ('MultiMask') Field Encoding	15 - 19
Table 15.33	Enable Receiver ('RxEn') Field Encoding	15 - 19
Table 15.34	Serial LSB Mode Register ('SerialLSBModeReg') Bit Assignments	15 - 20
Table 15.35	Clock Sampling Mode ('ClkSample') Field Encoding	15 - 20
Table 15.36	Sync Character Mode ('SyncChar') Field Encoding	15 - 20
Table 15.37	Stop Bit Mode ('StopBit') Field Encoding	15 - 20
Table 15.38	Parity Mode ('ParityMode') Field Encoding	15 - 20
Table 15.39	Parity Enable ('ParityEn') Field Encoding	15 - 21
Table 15.40	Serial Transmitter Control Register ('SerialTxControl Reg(1:0)') Bit Assignments	15 - 21
Table 15.41	DTR Status ('DTRN') Field Encoding	15 - 21
Table 15.42	Transmitter Bits per Character ('TxBitsPerChar')	15 - 22
Table 15.43	Transmitter Break ('TxBreak') Field Encoding	15 - 22
Table 15.44	Transmitter Enable ('TxEn') Field Encoding	15 - 22
Table 15.45	CRC Type ('CRCType') Field Encoding	15 - 22
Table 15.46	RTS Status ('RTSN') Field Encoding	15 - 22
Table 15.47	Transmitter CRC Enable ('TxCRCEn') Field Encoding	15 - 22
Table 15.48	Serial Reset Register ('SerialResetReg') Bit Assign- ments	15 - 24

Table 15.49	Master Reset ('Reset') Field Encoding	15 - 24
Table 15.50	Interrupt Enable ('IntEn') Field Encoding	15 - 24
Table 15.51	Serial MSB Mode Register ('SerialModeReg') Bit Assignments	15 - 25
Table 15.52	CRC Preload ('CRCPreload') Field Encoding	15 - 25
Table 15.53	Endec Mode ('EndecMode') Field Encoding	15 - 25
Table 15.54	Automatic SDLC Poll Frame Mode ('AutoPollFrame') Field Encoding	15 - 25
Table 15.55	Automatic Mark or Flag during Transmitter Idle ('IdleMark') Field Encoding	15 - 25
Table 15.56	Automatic Transmitter Underrun Abort Field Encoding	15 - 26
Table 15.57	SDLC Loop Mode ('SDCLoop') Field Encoding	15 - 26
Table 15.58	Sync Char Length ('SyncLength') Field Encoding	15 - 26
Table 15.59	Serial LSB Clock Control Register ('SerialSBClkControlReg') Bit Assignments	15 - 26
Table 15.60	Receiver Clock Input Source ('RxClkIn') Field Encoding	15 - 27
Table 15.61	Transmitter Clock Input Source ('TxClkIn') Field Encoding	15 - 27
Table 15.62	SerialSecondaryClk() Direction ('SecClkDir') Field Encoding	15 - 27
Table 15.63	SerialSecondaryClk() Output Clock Select ('SecClkOutSel') Field Encoding (if output is enabled)	15 - 27
Table 15.64	PLL Control Register ('PLLControlReg')	15 - 29
Table 15.65	PLL Mode ('PLLMode') Field Encoding	15 - 29
Table 15.66	Loopback Test Enable ('Loopback') Field Encoding	15 - 29
Table 15.67	Echo Test Enable ('Echo') Field Encoding	15 - 30
Table 15.68	Baud Rate Clock Generator Clock Input ('BRCGClkIn') Field Encoding	15 - 30
Table 15.69	Baud Rate Clock Generator Enable ('BRCGEn') Field Encoding	15 - 30
Table 15.70	Serial MSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments	15 - 31
Table 15.71	Interrupt Enable Bit Field Encoding	15 - 31
Table 15.72	MetaMeta Control Register Enable ('MetaMetaRegEn')	15 - 31
Table 15.73	Serial LSB and MSB Frame Status FIFO Data Register ('SerialMSBFrameStatusDataReg' and ('SerialMSBFrameStatusDataReg') Bit Assignments	15 - 32
Table 15.74	Frame Status FIFO Overflow Flag ('Overflow')	15 - 32
Table 15.75	Frame Status Not Empty Flag ('NotEmpty')	15 - 32
Table 15.76	Serial Meta Meta Register ('SerialMetaMetaReg') Bit Assignments	15 - 33
Table 15.77	Wait for Receiver CRC ('WaitCRC') Field Encoding	15 - 33
Table 15.78	Automatically Drive SerialTxData High when disabled ('AutoTxDataHigh') Field Encoding	15 - 33
Table 15.79	Automatically De-assert SerialRTS after SDLC Flag ('AutoSDLCRTS') Field Encoding	15 - 34
Table 15.80	Automatically Reset Transmitter Underrun ('AutoTxUnderrun') Field Encoding	15 - 35
Table 15.81	Automatic Insertion of Transmitter Opening Flag ('AutoTxFlag') Field Encoding	15 - 35
Table 16.1	Compatible Forward Data Transfer Variations	16 - 3
Table 16.2	Interrupt Descriptions	16 - 6
Table 16.3	Bidirectional Parallel Port Interface Centronics Controller Registers	16 - 9
Table 16.4	Centronics Sub ModeRegister ('CentSubModeReg') Bit Assignments	16 - 9
Table 16.5	Centronics Compatible Sub Mode ('SubMode') Field Encoding	16 - 9

Table 16.6	Centronics Status Register ('CentStatusReg') Bit Assignments	16 - 10
Table 16.7	IoCS(7:6) Mask Enable Field Encoding	16 - 10
Table 16.8	Centronics Negotiation Interrupt Pending Field	16 - 10
Table 16.9	Centronics Negotiation Interrupt Enable Field	16 - 11
Table 16.10	Centronics Reset Interrupt Pending Field	16 - 11
Table 16.11	Centronics Negotiation Interrupt Enable Field	16 - 11
Table 16.12	Printer Error Field ('PErr') Field Encoding	16 - 11
Table 16.13	Select ('Select') Field Encoding	16 - 11
Table 16.14	Printer Fault ('Fault') Field Encoding	16 - 11
Table 16.15	Printer Acknowledge Negated ('AckN') Field Encoding	16 - 12
Table 16.16	Printer Busy Field ('Busy') Field Encoding	16 - 12
Table 16.17	Centronics Control Register ('CentControl') Bit Assignments	16 - 12
Table 16.18	Negotiation XFlag Reply ('NegRep') Field Encoding	16 - 12
Table 16.19	Negotiation Mode ('NMode') Field Encoding	16 - 13
Table 16.20	Centronics Nibble Data Register ('CentNibbleDataReg') Bit Assignments	16 - 13
Table 16.21	Centronics Host Status Register ('CentHostStatusReg') Bit Assignments	16 - 14
Table 16.22	AutoFeed Negated ('nAutoFeed') Field Encoding	16 - 14
Table 16.23	Initialize Negated ('nInit') Field Encoding	16 - 14
Table 16.24	Select In Negated ('nSelectIn') Field Encoding	16 - 14
Table 16.25	Host Strobe Negated ('nStrobe') Field Encoding	16 - 14
Table 16.26	Centronics Minimum Delay Register ('CentDelayReg') Bit Assignments	16 - 15
Table 16.27	2500ns Delay Type Field ('D2500ns') Field Encoding	16 - 15
Table 16.28	500ns Delay Type Field ('D500ns') Field Encoding	16 - 15
Table 16.29	Example Settings for Delay Type Fields	16 - 15
Table 17.1	Laser Printer Video Controller Registers	17 - 5
Table 17.2	Laser Video Control Register ('LaserControlReg') Bit Assignments	17 - 5
Table 17.3	FIFO Full Field	17 - 5
Table 17.4	Request Size Field	17 - 6
Table 17.5	Reset FIFO Field	17 - 6
Table 17.6	Last Band ('Last') Field Encoding	17 - 6
Table 17.7	Video Direction ('Dir') Field Encoding	17 - 6
Table 17.8	Inverse Video Data ('Inv') Field Encoding	17 - 6
Table 17.9	Phase Lock Loop Enabled ('PLL') Field Encoding	17 - 7
Table 17.10	Laser Video Vertical Skip Register ('LaserVertSkipReg') Bit Assignments	17 - 7
Table 17.11	Vertical Skip Enable ('VSDis') Field Encoding	17 - 7
Table 17.12	Vertical Skip Count ('VSCount') Field Encoding	17 - 7
Table 17.13	Laser Video Horizontal Skip Register ('LaserHorSkipReg') Bit Assignments	17 - 8
Table 17.14	Horizontal Skip Enable ('HSDisEn') Field Encoding	17 - 8
Table 17.15	Horizontal Skip Count ('HSCount') Field Encoding	17 - 8
Table 17.16	Laser Video Band Line Count Register ('LaserBandLineCountReg') Bit Assignments	17 - 9
Table 17.17	Band Line (Lines per Band/Page) Count ('BLCount') Field Encoding	17 - 9
Table 17.18	Laser Video Horizontal Pixel Word Count Register ('LaserPixelWordCountReg') Bit Assignments	17 - 9
Table 17.19	Video Vertical Pixel Count ('PixCount') Field Encoding	17 - 10
Table 18.1	R36100 Reset Configuration Mode Features	18 - 1
Table 18.2	Boot Prom Reset Configuration Modes for ExcSIntN(2:1) Pins	18 - 2
Table 19.1	Reserved Emulator Addresses	19 - 2
Table 19.2	Debug Interface Register Address Assignments	19 - 4

Table 19.3	Debug Tracepoint Address Register ('DebugTraceAddr Reg') Bit Assignments	19 - 5
Table 19.4	Table Debug Tracepoint Control Register ('DebugTrace ControlReg') Bit Assignments	19 - 5
Table 19.5	Cause is a Tracepoint ('CTP') Field Encoding	19 - 5
Table 19.6	Tracepoint Enable ('TP') Field Encoding	19 - 6
Table 19.7	Debug Control Register ('DebugControlReg') Bit Assign- ments	19 - 6
Table 19.8	Writability ('Wr') Field Encoding	19 - 6



List of Figures

Figure No.	Figure Title	Page
Figure 1.1	R36100 Block Diagram.....	1 - 1
Figure 1.2	R36100 Logic Symbol.....	1 - 9
Figure 1.3	R36100-based Printer System.....	1 - 14
Figure 1.4	Development Support	1 - 15
Figure 2.1	CPU Registers.....	2 - 2
Figure 2.2	Instruction Encoding	2 - 3
Figure 2.3	Byte Ordering Conventions	2 - 5
Figure 2.4	Unaligned Words	2 - 6
Figure 2.5	5-Stage Pipeline	2 - 9
Figure 2.6	5-Instructions per Clock Cycle	2 - 10
Figure 2.7	Lead Delay	2 - 11
Figure 2.8	Branch Delay	2 - 11
Figure 3.1	Cache Line Selection	3 - 2
Figure 3.2	R36100 Execution Core and Cache Interface	3 - 5
Figure 3.3	Phased Access of Instruction and Data Caches	3 - 6
Figure 3.4	R36100 Instruction Cache Index Address Path	3 - 10
Figure 3.5	R36100 Cache in One Portion	3 - 12
Figure 3.6	R36100 Cache in Two Portions	3 - 12
Figure 3.7	R36100 Cache in Four Portions	3 - 12
Figure 4.1	Virtual Address Format	4 - 2
Figure 4.2	Virtual-to-Physical Address Translation in R36100	4 - 4
Figure 5.1	R36100 CPO Registers	5 - 1
Figure 5.2	R36100 Cache Control Register	5 - 2
Figure 5.3	R36100 Cause Register	5 - 7
Figure 5.4	R36100 Status Register	5 - 9
Figure 5.5	R36100 PrID Register	5 - 11
Figure 6.1	R36100 Cause Register	6 - 4
Figure 6.2	The Status Register	6 - 6
Figure 6.3	Pipelining in the R3051 Family	6 - 7
Figure 6.4	Synchronized Interrupt Operation	6 - 10
Figure 6.5	Direct Interrupt Operation	6 - 11
Figure 6.6	Synchronized BrCond Inputs	6 - 12
Figure 6.7	Kernel and Interrupt Status Being Saved on Interrupts	6 - 13
Figure 6.8	Code Sequence to Initialize Exception Vectors	6 - 14
Figure 6.9	Preserving Processor Context	6 - 15
Figure 6.10	Exception Cause Decoding	6 - 16
Figure 6.11	Exception Service Branch Table	6 - 17
Figure 6.12	Returning from Exception	6 - 18
Figure 6.13	Polling System Using BrCond	6 - 20
Figure 6.14	Using BrCond for Fast Interrupt Decoding	6 - 21
Figure 7.1	R36100 Bus Interface Unit Block Diagram	7 - 1
Figure 7.2	CPU Latency to Start of Read	7 - 10
Figure 7.3	Start of Bus Read Operation	7 - 11
Figure 7.4	Data Sampling	7 - 12
Figure 7.5	Read Cycle Termination	7 - 14
Figure 7.6	Internal Processor States on 4-word Burst Read	7 - 15
Figure 7.7	Instruction Streaming Internal Operation Example	7 - 16
Figure 7.8	Write Cycle Termination	7 - 21
Figure 7.9	Write-Buffer-Full Operation	7 - 22
Figure 8.1	R36100 Memory Bus Controller Block Diagram	8 - 2
Figure 8.2	Memory and I/O MSB Base Address Register (‘MemMSBBaseAddrReg’)	8 - 9

Figure 8.3	Memory and I/O LSB Base Address Register ('MemLSB BaseAddrReg')	8 - 9
Figure 8.4	Memory and I/O MSB Page Mask Register ('MemMSB PageMaskReg')	8 - 11
Figure 8.5	Memory and I/O LSB Page Mask Address Register ('MemLSBPageMaskReg')	8 - 11
Figure 8.6	Memory and I/O Control Register Bit Assignments	8 - 12
Figure 8.7	Memory LSB Wait State Register ('MemLSBWaitState Reg')	8 - 14
Figure 8.8	Memory MSB Wait State Register ('MemMSBWaitState Reg')	8 - 16
Figure 8.9	1-Datum Read with 0 Wait-States	8 - 19
Figure 8.10	1-Datum Read with 0 Wait-States Using an Odd Chip Select	8 - 20
Figure 8.11	1-Datum Read with 1 Wait-State using StartRepeat Field	8 - 21
Figure 8.12	Read with Wait-State using RdStart2Datum Field	8 - 22
Figure 8.13	Read with Wait-State using SysWait	8 - 23
Figure 8.14	4-Word Burst Read with 0 Wait-States	8 - 24
Figure 8.15	4-Word Burst Read with Wait-States using RdDatum 2Datum Field	8 - 25
Figure 8.16	PCMCIA-Style Memory Read with 0 Wait-States	8 - 26
Figure 8.17	1-Datum Write with 0 Wait-States	8 - 28
Figure 8.18	1-Datum Write with 0 Wait-States using FCT245-Type Field	8 - 29
Figure 8.19	1-Datum Write with Wait-State using StartRepeat Field	8 - 30
Figure 8.20	1-Datum Write with Wait-State using WrStart2Datum Field	8 - 31
Figure 8.21	1-Datum Write with Wait-State using SysWait	8 - 32
Figure 8.22	Multi-Datum Burst Write	8 - 33
Figure 8.23	Multi-Datum Burst Write using Wait-State with WrDatum2Datum	8 - 34
Figure 8.24	PCMCIA-Style Memory Write with 0 Wait-States	8 - 35
Figure 8.25	Interleaved Read using FCT260-Type Field	8 - 37
Figure 8.26	Interleaved "Even" Read of FCT260-Type Memory	8 - 38
Figure 8.27	Interleaved "Odd" Read of FCT260-Type Memory	8 - 39
Figure 8.28	Interleaved Read using FCT245-Type Field	8 - 40
Figure 8.29	Interleaved Read using FCT543-Type Field	8 - 41
Figure 8.30	"Even" Read of FCT543-Type Memory	8 - 42
Figure 8.31	"Odd" Read of FCT543-Type Memory	8 - 43
Figure 8.32	Interleaved Write using FCT260-Type and FCT543-Type Fields	8 - 44
Figure 8.33	Interleaved Write using FCT245-Type Field	8 - 45
Figure 8.34	32-bit SRAM System	8 - 46
Figure 8.35	32-bit SRAM System using an Odd Chip Select	8 - 46
Figure 8.36	32-bit SRAM System using FCT245-Type	8 - 47
Figure 8.37	Interleaved FCT245-Type System	8 - 48
Figure 8.38	Interleaved FCT260-Type System	8 - 49
Figure 8.39	Interleaved FCT543-Type System	8 - 50
Figure 8.40	16-bit Big Endian SRAM System	8 - 51
Figure 8.41	16-bit Little Endian SRAM System	8 - 51
Figure 8.42	8-bit Big Endian SRAM System	8 - 52
Figure 8.43	8-bit Little Endian SRAM System	8 - 52
Figure 9.1	R36100 I/O Bus Controller Block Diagram	9 - 2
Figure 9.2	Memory and I/O Control Register ('MemIoCntrl Reg(7..0)')	9 - 9
Figure 9.3	I-Type I/O Read with 0 Wait-States	9 - 12
Figure 9.4	M-Type I/O Read with 0 Wait-States	9 - 13
Figure 9.5	PCMCIA-Style I/O Read with 0 Wait-States	9 - 14
Figure 9.6	I-Type I/O Write with 0 Wait-States	9 - 15
Figure 9.7	M-Type I/O Write with 0 Wait-States	9 - 16

Figure 9.8	PCMCIA-Style I/O Write with 0 Wait-States	9 - 17
Figure 9.9	I/O Read with Internal Wait-States	9 - 18
Figure 9.10	I/O Read with external SysWait Wait-State	9 - 19
Figure 9.11	I/O Write with Internal Wait-States	9 - 20
Figure 9.12	I/O Write with external SysWait Wait-State	9 - 21
Figure 9.13	I-Type I/O System with Direct Bus Connection	9 - 22
Figure 9.14	M-Type I/O System with Direct Bus Connection	9 - 22
Figure 9.15	I-Type I/O System using FCT245 Transceivers	9 - 23
Figure 9.16	M-Type I/O System using FCT245 Transceivers	9 - 23
Figure 9.17	I-Type I/O System using FCT543 Transceivers	9 - 24
Figure 9.18	M-Type I/O System using FCT543 Transceivers	9 - 24
Figure 9.19	16-bit I/O System with Big Endian Connection	9 - 25
Figure 9.20	16-bit I/O System with Little Endian Connection	9 - 25
Figure 9.21	8-bit I/O System with Little Endian Connection	9 - 26
Figure 9.22	8-bit I/O System with Little Endian Connection	9 - 26
Figure 10.1	R36100 DRAM Bus Controller Block Diagram	10 - 2
Figure 10.2	DRAM Refresh Count Register ('DramRefreshCount Reg')	10 - 9
Figure 10.3	DRAM Refresh Compare Register	10 - 10
Figure 10.4	DRAM RAS Mux Select Register ('DramRasMux SelReg')	10 - 11
Figure 10.5	DRAM CAS Mux Select Register	10 - 11
Figure 10.6	DRAM MSB Base Address Register ('DramMSBBase AddrReg')	10 - 14
Figure 10.7	DRAM MSB Bank Mask ('DramMSBBankMask(3:0)' Registers)	10 - 15
Figure 10.8	DRAM LSB Bank Control Register ('DramLSBControl Reg')	10 - 15
Figure 10.9	DRAM MSB Bank Control Register ('DramMSBControl Reg')	10 - 18
Figure 10.10	Basic DRAM Read	10 - 23
Figure 10.11	RAS Asserted at End of Transfer	10 - 24
Figure 10.12	RAS Asserted at Start of Transfer	10 - 25
Figure 10.13	RAS Precharge at Start of Transfer	10 - 26
Figure 10.14	Extended Row Address Hold	10 - 27
Figure 10.15	Extended Address Set-up	10 - 28
Figure 10.16	Extended CAS Width	10 - 29
Figure 10.17	Multiple Datum Read	10 - 30
Figure 10.18	Basic DRAM Write	10 - 32
Figure 10.19	RAS Asserted Throughout DRAM Write	10 - 34
Figure 10.20	Write Bus Turn-around	10 - 35
Figure 10.21	Two Datum Write	10 - 36
Figure 10.22	Interleaved FCT245-Type Read	10 - 37
Figure 10.23	Interleaved FCT260-Type Read	10 - 38
Figure 10.24	Single Word Access to Even Bank of FCT260-Type System	10 - 39
Figure 10.25	Single Word Access to Odd Bank of FCT260-Type System	10 - 40
Figure 10.26	Interleaved FCT543 Read	10 - 41
Figure 10.27	Single Word Access to Even Bank of FCT543-Type System	10 - 42
Figure 10.28	Single Word Access to Odd Bank of FCT543-Type System	10 - 43
Figure 10.29	Interleaved FCT245-Type Writes	10 - 45
Figure 10.30	Interleaved FCT260, FCT543-Type Writes	10 - 47
Figure 10.31	DRAM Staggered Refresh	10 - 48
Figure 10.32	Interleaved FCT245 Interface	10 - 50
Figure 10.33	Interleaved FCT260 Interface	10 - 51
Figure 10.34	Interleaved FCT543 Interface	10 - 52
Figure 11.1	DMA Controller Address and Data Flow Diagram	11 - 1
Figure 11.2	Rotating Priority Scheme	11 - 3
Figure 11.3	Internal DMA Algorithm	11 - 5

Figure 11.4	Internal DMA LSB Source Address Register ('DmaLSBSourceAddrReg')	11 - 10
Figure 11.5	Internal DMA MSB Source Address Register ('DmaMSBSourceAddrReg')	11 - 10
Figure 11.6	Internal DMA LSB Target Address Register ('DmaLSBTargetAddrReg')	11 - 10
Figure 11.7	Internal DMA MSB Target Address Register ('DmaMSBTargetAddrReg')	11 - 11
Figure 11.8	DMA LSB Count Register ('DmaLSBCountReg')	11 - 11
Figure 11.9	Internal DMA MSB Count Register ('DmaMSBCountReg')	11 - 11
Figure 11.10	Internal DMA LSB Control Register ('DmaLSBControlReg')	11 - 12
Figure 11.11	Internal DMA MSB Control Register ('DmaMSBControlReg')	11 - 15
Figure 11.12	External DMA Control Register ('ExtDmaControlReg') ..	11 - 17
Figure 11.13	External DMA Single Data Read using the Memory Controller (Data Transfer from Memory to Device)	11 - 19
Figure 11.14	External DMA Single Data Write using the Memory Controller (Data Transfer from Device to Memory)	11 - 20
Figure 11.15	External DMA Two-Data Burst Read using the Memory Controller (Data Transfer from Memory to Device)	11 - 21
Figure 11.16	External DMA Two-Data Burst Write using the Memory Controller (Data Transfer from Device to Memory)	11 - 21
Figure 12.1	PIO Block Diagram	12 - 1
Figure 12.2	PIO Data Register ('PioDataReg')	12 - 3
Figure 12.3	PIO Direction Register ('PioDirReg')	12 - 3
Figure 12.4	PIO Effect Select Register ('PioEffectSelReg')	12 - 4
Figure 13.1	Expansion Interrupt Controller (to CPU Interrupt)	13 - 1
Figure 13.2	Expansion Interrupt Controller: Steering Interrupts to DMA Requests	13 - 2
Figure 13.3	Expansion Interrupt Mask Register ('ExpIntMaskReg')	13 - 3
Figure 13.4	Expansion Interrupt Pending Register ('ExpIntPendReg')	13 - 3
Figure 13.5	Expansion Interrupt DMA Select Register ('ExpIntDMASelectReg')	13 - 5
Figure 14.1	Block Diagram of the R36100 Timers	14 - 1
Figure 14.2	Timer Prescaler Count Register ('TimerPrescalerCountReg')	14 - 3
Figure 14.3	Timer Count Register ('TimerCountReg')	14 - 3
Figure 14.4	Timer Compare Register ('TimerCompareReg')	14 - 4
Figure 14.5	Timer Pulse Width Modulation Register ('TimerPWMReg')	14 - 4
Figure 14.6	Timer Control Register ('TimerControlReg')	14 - 4
Figure 15.1	Serial Channel Interface Block Diagram	15 - 1
Figure 15.2	Serial LSB Status Register ('SerialLSBStatusReg')	15 - 9
Figure 15.3	Serial MSB Status Register ('SerialMSBStatusReg')	15 - 10
Figure 15.4	Serial Interrupt Vector ('SerialIntVector(1)')	15 - 10
Figure 15.5	Serial MSB Sync Character/Address Register ('SerialMSBSyncCharReg')	15 - 12
Figure 15.6	Serial LSB Sync Character/Address Register ('SerialLSBSyncCharReg')	15 - 12
Figure 15.7	Receiver Data FIFO Register ('SerialDataFIFOReg')	15 - 12
Figure 15.8	Serial Clock Status Register ('SerialClkStatusReg')	15 - 12
Figure 15.9	Serial MSB Baud Rate Clock Generator Compare Register ('SerialMSBBaudCompareReg')	15 - 13
Figure 15.10	Serial LSB Baud Rate Clock Generator Compare Register ('SerialLSBBaudCompareReg')	15 - 13

Figure 15.11	Serial MSB Interrupt Enable Register ('SerialIntEn Reg')	15 - 14
Figure 15.12	Serial Meta Register ('SerialMetaRegister')	15 - 14
Figure 15.13	Serial LSB Interrupt Enable Register ('SerialIntEn Reg')	15 - 15
Figure 15.14	Serial Interrupt Vector ('SerialIntVector(1)')	15 - 17
Figure 15.15	Serial Receiver Control Register ('SerialRxControl Reg')	15 - 17
Figure 15.16	Serial LSB Mode Register ('SerialLSBMode Reg')	15 - 19
Figure 15.17	Serial Transmitter Control Register ('SerialTxControl Reg')	15 - 21
Figure 15.18	Serial MSB Sync Character/Address Register ('Serial MSBSyncCharReg')	15 - 23
Figure 15.19	Serial LSB Sync Character/Address Register ('Serial LSBSyncCharReg')	15 - 23
Figure 15.20	Transmitter Data FIFO Register ('SerialDataFIFO Reg')	15 - 24
Figure 15.21	Serial MSB Mode Register ('SerialModeReg')	15 - 24
Figure 15.22	Serial LSB Clock Control Register ('SerialLSBClk ControlReg')	15 - 26
Figure 15.23	Serial MSB Baud Rate Clock Generator Compare Register ('SerialMSBBaudCompareReg')	15 - 28
Figure 15.24	Serial LSB Baud Rate Clock Generator Compare Register ('SerialLSBBaudCompareReg')	15 - 28
Figure 15.25	PLL Control Register ('PLLControlReg')	15 - 29
Figure 15.26	Serial MSB Interrupt Enable Register ('SerialIntEn Reg')	15 - 30
Figure 15.27	Serial MSB Frame Status FIFO Data Register ('SerialMSBFrameStatusDataReg')	15 - 32
Figure 15.28	Serial LSB Frame Status FIFO Data Register ('SerialLSBFrameStatusDataReg')	15 - 32
Figure 15.29	Serial Meta Meta Register ('SerialMetaMeta Reg')	15 - 33
Figure 16.1	Block Diagram of the Bidirectional Parallel Port	16 - 1
Figure 16.2	Centronics Sub Mode Register ('CentSubMode Reg')	16 - 9
Figure 16.3	Centronics Status Register ('CentStatusReg')	16 - 9
Figure 16.4	Centronics Control Register ('CentControl Reg')	16 - 12
Figure 16.5	Centronics Nibble Data Register ('CentNibbleData Reg')	16 - 13
Figure 16.6	Centronics Host Status Register ('CentHostStatus Reg')	16 - 13
Figure 16.7	Centronics Minimum Delay Register ('CentDelay Reg')	16 - 14
Figure 16.8	Typically Classic Compatible Mode Transaction	16 - 16
Figure 16.9	Typical Parallel Port System Connections	16 - 16
Figure 17.1	Block Diagram of the Laser Printer Video Port	17 - 1
Figure 17.2	Laser Video Control Register ('LaserControl Reg')	17 - 5
Figure 17.3	Laser Vertical Skip Register ('LaserVertSkip Reg')	17 - 7
Figure 17.4	Laser Video Horizontal Skip Register ('LaserHorSkip Reg')	17 - 8
Figure 17.5	Laser Video Band Line Count Register ('LaserBandLine CountReg')	17 - 9
Figure 17.6	Laser Video Horizontal Word Count Register ('Laser HorizPixelCountReg')	17 - 9
Figure 17.7	Laser Video Data Register ('LaserWordDataReg')	17 - 10
Figure 17.8	R36100 Video Interface System Example	17 - 10

Figure 18.1	Cold Start	18 - 4
Figure 18.2	Warm Reset	18 - 4
Figure 18.3	Configuration Mode Initialization Logic	18 - 5
Figure 18.4	Mode Vector Timing	18 - 5
Figure 18.5	Reset Timing	18 - 6
Figure 18.6	R36100 Clocking	18 - 6
Figure 19.1	Debug Tracepoint Address Register ('MSB Debug TraceAddrReg')	19 - 4
Figure 19.2	Debug Tracepoint Address Register ('LSB Debug TraceAddrReg')	19 - 4
Figure 19.3	Debug Tracepoint Control Register ('DebugTrace ControlReg')	19 - 5
Figure 19.4	Debug Control Register ('DebugControlReg')	19 - 6
Figure 19.5	R36100 SysClk Phase Initialization Case A	19 - 7
Figure 19.6	R36100 SysClk Phase Initialization Case B	19 - 7



Introduction

The IDT79R36100 is a highly integrated member of the IDT RISController family. The R36100 RISController incorporates the "system on a chip" integration philosophy and is well-suited for a wide variety of low-cost embedded applications.

The R36100 RISController contains the general purpose R3000A MIPS RISC CPU core and substantial amounts of on-chip Instruction Cache and Data Cache memory. In addition, the R36100 integrates four Memory Controllers on-chip, including ROM, DRAM, I/O, and DMA; printer and data communication peripherals, including an IEEE 1284 Parallel Port, Laser Printer Video Rasterizer, and two Serial Communications Ports; and standard embedded peripherals, including an Interrupt Controller, Timers, and Parallel Inputs and Outputs.

This extensive integration simplifies the overall system design and reduces external component requirements, system cost and development time.

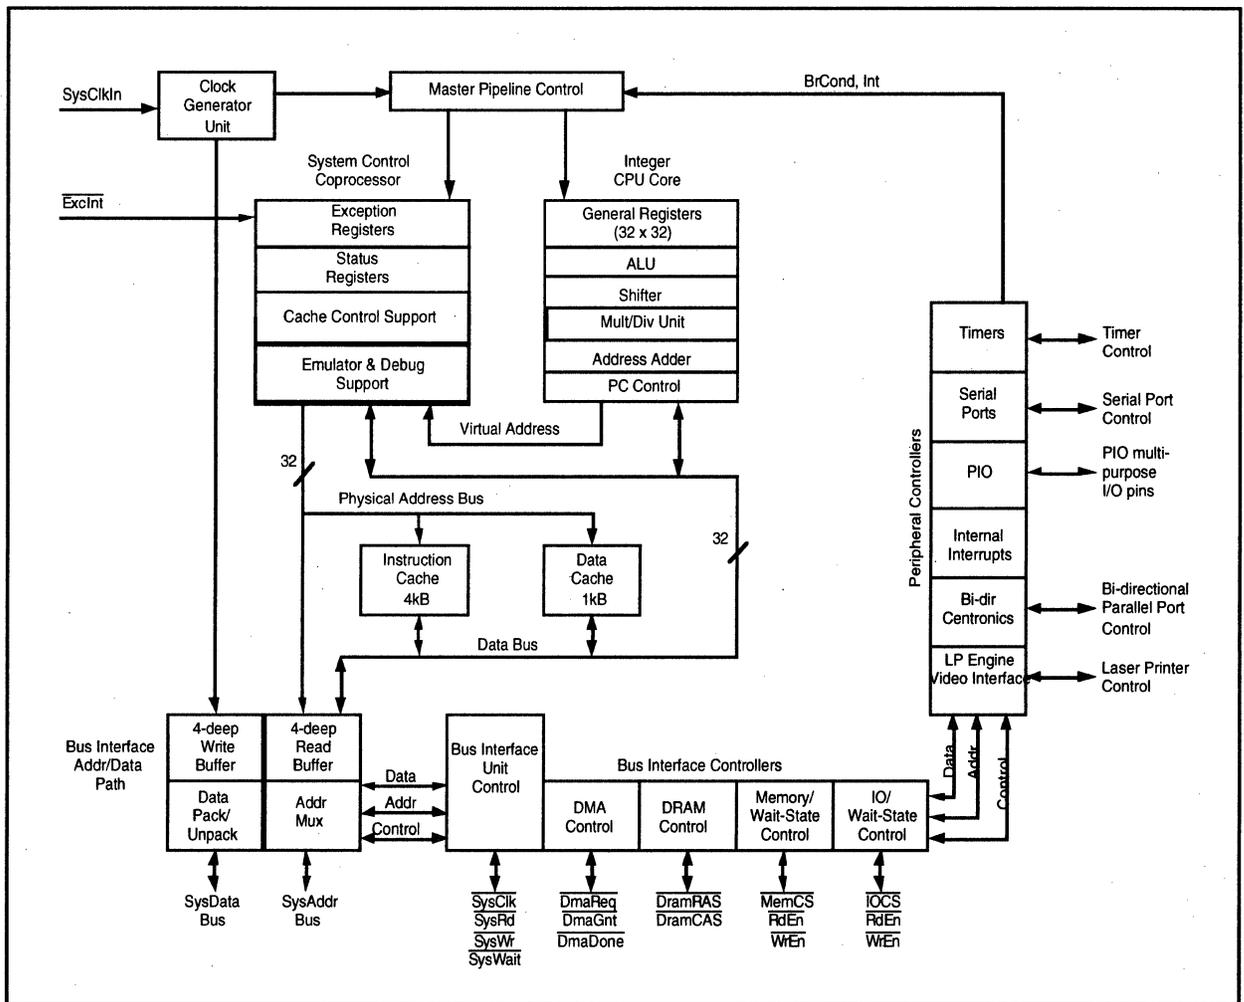


Figure 1.1 R36100 Block Diagram

The R36100 RISController is software compatible with all of the IDT RISController family, including the low cost 32-bit R3051 RISControllers and the R4xxx Orion family of high-performance 64-bit CPUs. Common instruction set architecture (ISA) enables the same applications software to be used across a wide variety of price/performance points.

The R36100 RISController integrates four on-chip Bus Controllers, allowing seamless interfacing with a wide variety of standard memories and peripherals that include:

- Standard page mode DRAMs
- EPROMs, FLASH, SRAM, Dual-Port SRAM
- FIFOs, SCSI, A/D, and other I/O peripherals
- Ethernet, Data Compression, and other coprocessors

The R36100 RISController integrates an IEEE Parallel Port, RS-232C and Local Talk Serial Ports, and a Laser Printer Video Rasterizer, to serve printer system applications that include:

- Monochrome laser and ink-jet printers
- Host based printer cards
- Multi-function laser/fax printer systems

The R36100 RISController integrates asynchronous and synchronous Serial Ports and multiple Timers, to serve data communications applications that include:

- Local Area Network (LAN) interface cards
- CSU/DSU SDLC/HDLC line driver cards
- Router, switcher, and data compressor cards

The next section contains a list of R36100 features.

R36100 Feature List

- Instruction set compatible with IDT79R3000A, R3051 family, and IDT Orion family of MIPS RISC CPUs
- System cost minimized through high level of integration
 - RISC CPU
 - Instruction Cache
 - Data Cache
 - Flexible Bus Interface
 - Controllers
 - Peripheral
- Double frequency clock input
- 24 MIPS/ 42K Dhrystones-2.1 at 25 MHz
- 3.3V and 5V versions
- Low cost PQFP packaging
- On-chip instruction and data caches
 - 4KB of Instruction Cache
 - 1KB of Data Cache
 - Improved Cache Control for fast data movement and cache locking
- Flexible bus interface allows simple, low cost designs
 - Separate de-multiplexed Address Bus and Data Bus
 - Synchronized Bus Interface Timing
 - On-chip 4-deep write buffer eliminates memory write stalls
 - On-chip 4-word read buffer supports burst or simple block reads
 - Programmable port width interface (8-,16-, and 32-bit memory sub-regions)
- On-chip DRAM Controller with Address Multiplexer
 - Supports non-interleaved or Interleaved DRAM memory
- On-chip Memory and I/O Controller
 - Chip Selects
 - Wait-State Generator
 - Supports non-interleaved or interleaved ROMs
 - Boot from 8-bit, 16-bit, 32-bit or interleaved ROMs
 - Supports CS/Rd/Wr I/O protocol
 - Supports CS/Wr/Strobe I/O protocol
 - Supports PCMCIA Master protocol
- On-chip DMA Controller for autonomous burst data movement
 - 4 internal channels
 - 2 external channels
 - On-chip Parallel I/O pins
 - On-chip Interrupt Expansion controller
- On-chip Timers
- On-chip Serial Port(s)
- On-chip IEEE 1284 Bidirectional Centronics Target Interface Controller
- On-chip Laser Printer Video Raster Engine Interface Controller
- "Reduced Frequency Mode" assists in power-managed and "Green PC" applications
- Complete software support
 - Optimizing compilers
 - Real-time operating systems
 - Monitors/debuggers
 - Floating Point emulation software
 - Printer Page Description Languages
- Built-in Debug/Emulator Support

Device Overview

The R36100 can be viewed as a “system on a chip,” the embodiment of a discrete system built around the R3000A CPU. Integrating system functions onto a single chip reduced the system’s cost, size, and power requirements. This high level of integration also reduced system complexity and minimized system development time.

Figure 1.1 on page 1 provides a block level representation of the R36100’s functional units. This section also includes the R36100’s logic symbol diagram (Figure 1.2 on page 9) and pin description table (Table 1.1 on page 10). A system overview is presented here in Chapter 1 with more detailed information provided in subsequent chapters.

CPU Core

The R36100 RISController is based on the R3000A CPU core. The R3000A is a full 32-bit RISC integer execution engine, capable of sustaining a peak single cycle execution rate by using its five-stage pipeline. The CPU core contains an integer ALU unit and bit shifter with a separate integer multiplier/divider unit, address adder and program counter generator, and 32 orthogonal 32-bit registers. The R36100 execution core implements the MIPS-I Instruction Set Architecture (ISA). Therefore, the R36100 is binary compatible with all other MIPS CPU engines, including the low cost R3051 family and the high-speed R4xxx Orion family.

System Control Co-Processor

The R36100 RISController integrates an on-chip System Control Co-processor (CPO). CPO manages the R36100’s exception handling operations, its virtual to physical address memory mapping, and its various programmable bus-to-cache interface capabilities. All of these topics are discussed in subsequent chapters.

The R36100 does not include the optional TLB found in other members of the IDT RISController family. Instead, the R36100 performs virtual to physical address mapping identical to that of the R3051 family’s Base Versions. These Base Version devices still support distinct kernel and user mode operation but do not require page management software or an on-chip TLB, leading to a simplified operating system software model and a lower cost processor.

Clock Generator Unit

The R36100 RISController is driven from a single, 2x-frequency input clock. An on-chip clock generator unit is responsible for managing the interaction of the CPU core, caches, and bus interface. The clock generator unit replaces the external delay line that was required in discrete R3000A based systems.

For power sensitive or “Green” applications, the R36100 supports a reduced frequency mode, allowing the system to reduce power consumption in idle periods.

Instruction Cache

The R36100 RISController integrates 4kB of on-chip Instruction Cache, organized with a line size of 16 bytes (four 32-bit entries). This relatively large cache contributes substantially to the high performance inherent in the R36100, which allows systems based on the R36100 to achieve high performance even from low-cost memory systems. The cache is implemented as a direct mapped cache and is capable of caching instructions from anywhere within the large physical address space. The cache is implemented using physical addresses and physical tags (rather than virtual addresses or tags), which does not require flushing on context switches.

The R36100 implements special features that allow the instruction cache to be split into halves or quarters; each section then services a different area of the large address space. This feature enables the system software to “lock” time critical code—such as router address hash-table lookup algorithms and interrupt service routines—into one of the halves or quarters while allowing other tasks to utilize unused areas without disrupting the time critical code. This technique permits software to perform instruction cache “locking” without requiring memory management support.

Data Cache

The R36100 RISController incorporates an on-chip data cache of 1KB organized as a line size of 4 bytes (one word). This relatively large data cache contributes substantially to the high performance of the R36100. As with the instruction cache, the data cache is implemented as a direct mapped physical address cache and is capable of mapping any word within the large physical address space.

The data cache is implemented as a write-through cache, to ensure that main memory is always consistent and coherent with the internal cache. To minimize processor stalls due to data write operations, the bus interface unit incorporates a 4-deep write buffer which captures address and data at the processor execution rate, allowing the data to be retired to main memory at a much slower rate without impacting the performance of the CPU core.

The R36100 contains special features that also allow the data cache to be split into halves or quarters; each section services a different area of the large address space. This feature enables the system software to “lock” time critical data—such as routing address information tables and the interrupt stack—into one of the halves or quarters while allowing other tasks to utilize unused portions without disrupting the critical data. This technique permits software to perform data cache “locking” without requiring memory management support.

Bus Interface Unit

The R36100 RISController uses its large internal caches to provide the execution engine with most of its memory bandwidth requirements. The execution engine pipeline can then perform both 1 instruction fetch and 1 data load/store per clock cycle. And only on the rare occasion of a cache miss or on writes does the R36100 require its external bus interface; therefore, the R36100 is able to use a simple bus interface that connects to slow, inexpensive memory devices.

The R36100 bus interface uses a de-multiplexed address and data bus. The bus interface readily connects to memory subsystems that are 8-, 16-, 32-bits wide, and/or interleaved 32-bit.

The R36100 incorporates a 4-deep write buffer to decouple the speed of the execution engine from the speed of the memory system. The write buffers capture and FIFO the processor's address and data information during internal store operations at the CPU pipeline rate. The write buffer then presents the bus interface write transactions at the rate the memory system can accommodate.

During main memory writes, the R36100 can break a large datum—such as a 32-bit word—into a series of smaller transactions—such as bytes—according to the width of the memory port being written. This operation is transparent to the software that initiated the store, ensuring that the same software is able to run in a variety of memory systems.

The R36100 read interface performs both single datum reads and quad word reads. To accommodate slower reads, the R36100 incorporates a 4-deep read buffer FIFO, allowing the external interface to queue up data within the processor before releasing it to perform a burst fill of the internal caches.

In addition, the R36100 can perform on-chip data packing when performing large datum reads—such as quad words—from narrower memory systems—such as 16-bits. Once again, this operation is transparent to the software, simplifying migration of software to different memory systems and simplifying field upgrades to wider memory. Since this capability works for either instruction or data reads, using 8-, 16-, 32-bit, or interleaved boot PROMs is easily supported by the R36100.

As described throughout this manual, one of the on-chip memory bus controllers services bus transactions. The bus interface unit merely provides a common translation between these memory bus controllers and the CPU core.

Memory Controller

The R36100 RISController uses the on-chip memory controller to gluelessly attach external ROM—including FLASH—and/or SRAM in a number of system configurations. For example, the memory controller supports interleaved ROM and/or SRAM, 8-bit boot ROM, 32-bit burst ROMs, as well as an array of simple 32-bit wide EPROMs. Under the control of boot software, the memory controller integrates all control signals and manages the access timing and wait-state generation for multiple banks.

DRAM Controller

The R36100 RISController integrates an on-chip DRAM controller. The DRAM controller directly controls up to four banks of standard page mode DRAMs in a number of configurations, including systems with varying densities of DRAM; 32-bit wide, interleaved DRAM; and 16-bit wide DRAM subsystems.

I/O Controller

To perform all necessary address decoding and wait-state generation for external I/O devices, the R36100 RISController has an on-chip I/O controller. In addition, the on-chip I/O controller interfaces as a master to PCMCIA, including support of the large address space required and the PCMCIA chip-select protocol and timing.

DMA Control and Interface

The R36100 RISController features on-chip DMA control for internal peripherals, external peripherals, and external memory. Multiple internal channels are provided, allowing block moves of data between any combination of memory and I/O device. Each channel can also be interrupt controlled so that an I/O peripheral—like the serial port—can regulate the individual transactions of a block move.

The R36100 RISController also supports external DMA masters that take over the external system bus via a bus request and grant handshake. Once in control, the external DMA master can read and write to memory, I/O, and internal peripherals via the R36100's bus controllers.

Counter/Timers

The R36100 RISController contains three general purpose timers. Each timer consists of a 16-bit count register as well as a 16-bit compare register. The count register resets to zero and then counts upward until it equals the compare register. When the count register equals the compare register, the TCN output is asserted and the count is reset back to zero.

To increase the amount of time each timer can handle, the timers use a common 16-bit prescaler counter. Each timer is programmable to select a power-of-2 divisor of the prescaler.

Using the default mode, each timer can be used as a general purpose real-time clock. Some special effects include:

- Bus timeout timer
- Watch dog timer
- PWM/square wave/baud rate generator
- Gated clock external event counter

PIO Interface

For controlling multi-purpose utility pins, the R36100 RISController has a Parallel Input/Output (PIO) interface. The PIO pins can be programmed to act as general purpose inputs or outputs.

Each PIO pin is multiplexed with other controller's inputs or outputs. This flexible arrangement allows the system designers to customize R36100's resources according to their needs. Therefore, designs needing a special purpose controller—such as the laser printer video controller—can allocate the LP Video pins for that purpose; other applications can use those pins for general purpose inputs or outputs.

Serial Communications Controller

The R36100 RISController integrates a dual channel serial port. This peripheral controller can perform a variety of synchronous and asynchronous protocols, including RS-232C, LocalTalk, SDLC, and HDLC. To maximize throughput, the on-chip Serial Port is optionally serviced by the auto-initiated on-chip DMA controller, which can automatically block move data to and from the port.

Interrupt Controller

The R36100 RISController integrates an on-chip interrupt controller to manage both external interrupts and interrupts signaled from the on-chip peripherals. The interrupt controller speeds interrupt service of the internal interrupts and assists in interrupt prioritization and nesting as well as interfacing with the auto-initiated DMA.

IEEE 1284 Bi-directional Centronics

The R36100 RISController includes an internal IEEE1284 parallel port peripheral, which implements a true bi-directional Centronics port. Features include:

- 8-bit input Target Compatible protocol (for backward compatibility with Centronics)
- Nibble and byte mode output protocol (for backward compatibility with PCs)
- ECP protocol (for the emerging Laser Printer PC standard)
- EPP protocol (for communications applications)
- External transceiver interface control pins
- Auto-initiated DMA via internal interrupts

Laser Printer Video Interface

The R36100 RISController integrates an on-chip laser printer video/control interface. This peripheral provides support for the following:

- 1-bit serial stream laser printer or raster engine interface
- On-chip FIFO
- Programmable margin widths and page lengths
- Auto-initiated DMA via internal interrupts

Logic Symbol

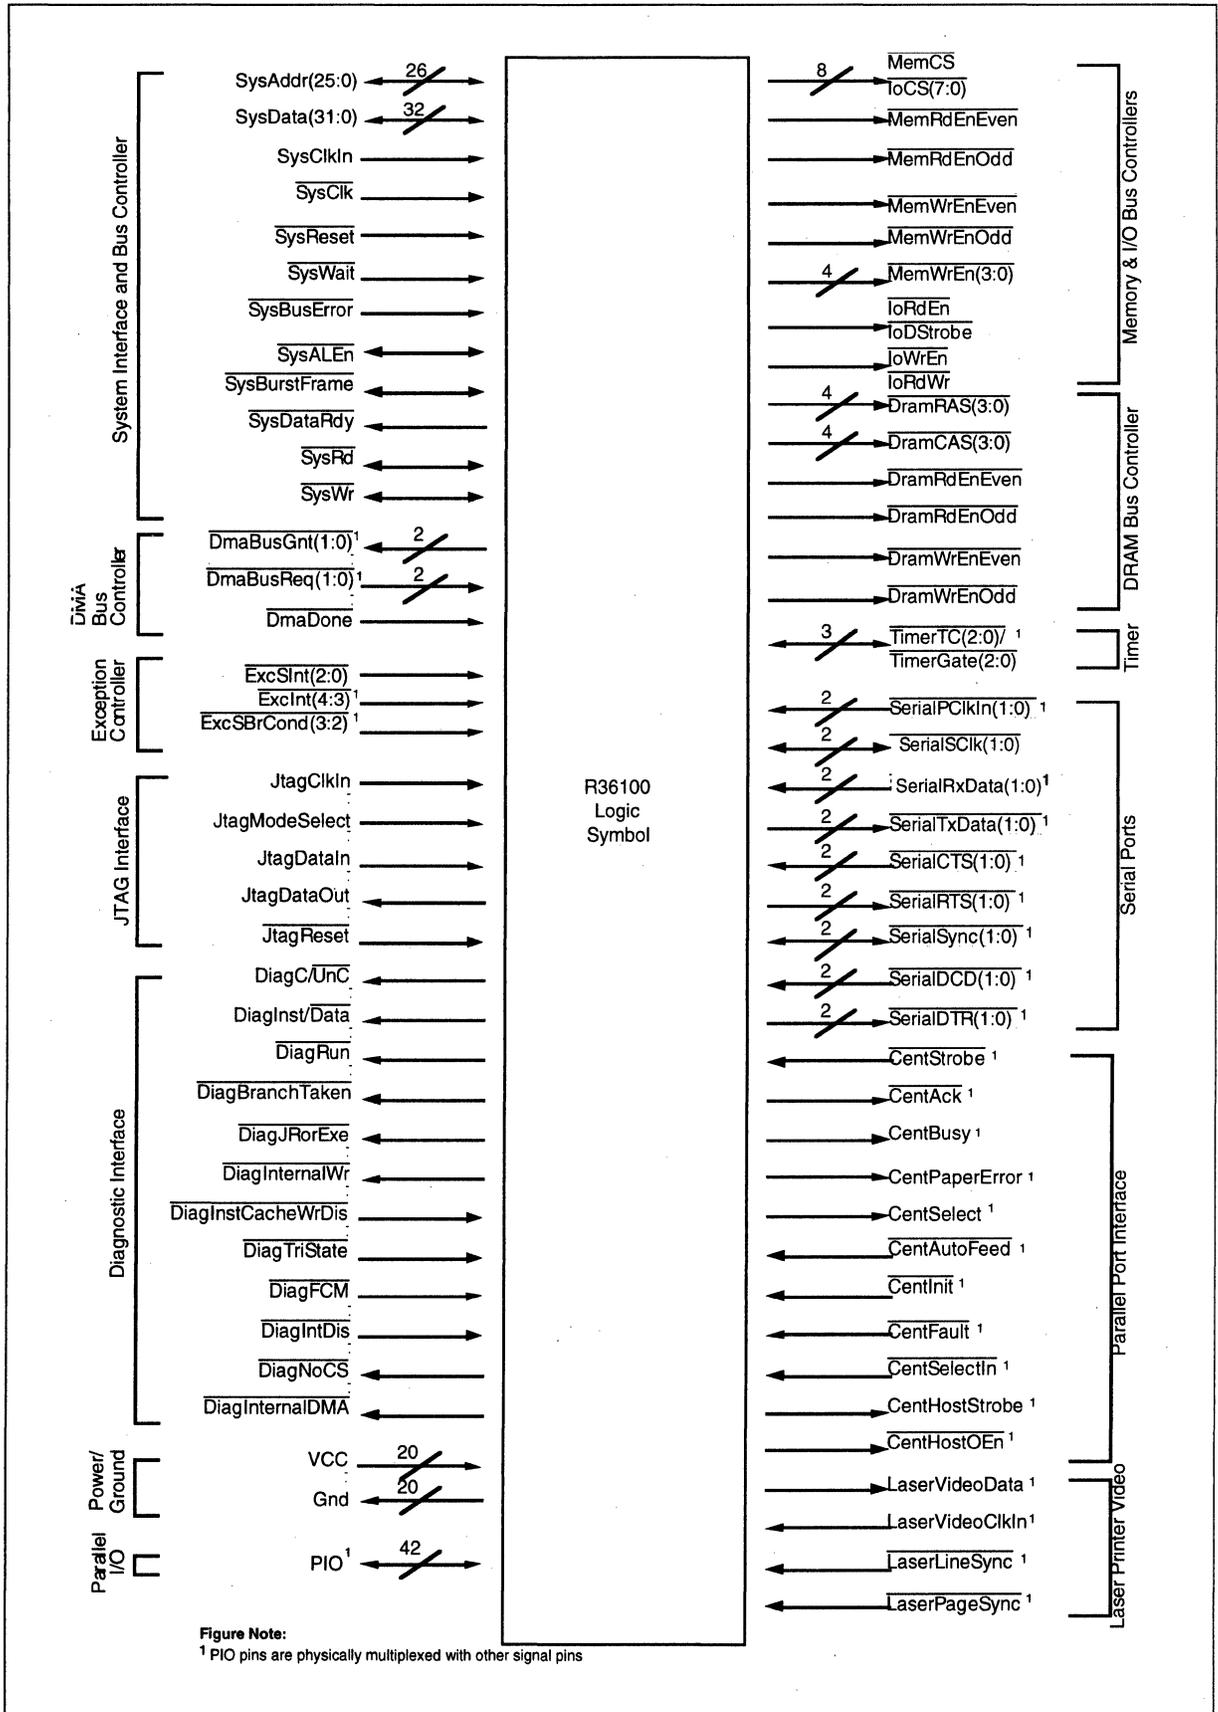


Figure 1.2 R36100 Logic Symbol

Pin Description

Pin Function and Description		
System Bus Interface Pins		
SysAddr(25:0)	I/O	System Address Bus. Also serves as the DramAddr(13:2) Bus.
SysData(31:0)	I/O	System Data Bus.
SysClkIn	Input	System Clock Input. Twice (2x) the internal CPU frequency.
$\overline{\text{SysClk}}$	Output	System Clock Output. All other outputs are referenced to this system clock.
$\overline{\text{SysReset}}$	Input	System Reset. Initializes entire chip, except for JTAG circuitry.
$\overline{\text{SysWait}}$	Input	System Wait. Extends current bus transaction.
$\overline{\text{SysBusError}}$	Input	System Bus Error. Terminates current bus transaction.
$\overline{\text{SysALEn}}$	I/O	System Address Latch Enable. Indicates valid address at the beginning of a bus transaction.
$\overline{\text{SysBurstFrame}}$	I/O	System Burst Frame. First indicates the beginning of a bus transaction. Then indicates if the bus transaction is a burst and if the next datum is the last datum.
$\overline{\text{SysDataRdy}}$	Output	System Data Ready. Indicates valid data during each datum of a bus transaction (except when $\overline{\text{SysWait}}$ is asserted).
$\overline{\text{SysRd}}$	I/O	System Read. Indicates current bus transaction is a read.
$\overline{\text{SysWr}}$	I/O	System Write. Indicates current bus transaction is a write.
DRAM Controller Pins		
$\overline{\text{DramRAS(3:0)}}$	Output	DRAM Row Address Strobe.
$\overline{\text{DramCAS(3:0)}}$	Output	DRAM Column Address Strobe.
$\overline{\text{DramRdEnEven}}$	Output	DRAM Read Enable for Even FCT245/543 Type Banks. On FCT260 type banks, it is the read enable for both even and odd banks.
$\overline{\text{DramRdEnOdd}}$	Output	DRAM Read Enable for Odd FCT245/543 Type Banks. On FCT260 Type banks, it is the path select.
$\overline{\text{DramWrEnEven}}$	Output	DRAM Write Enable for Even Banks.
$\overline{\text{DramWrEnOdd}}$	Output	DRAM Write Enable for Odd Banks.
Memory Controller Pins		
$\overline{\text{MemCS/IoCS(7:0)}}$	Output	Memory or I/O Chip Selects. MemCS(0) and optionally MemCS(1) are reserved for the Boot PROM. IoCS(6) and/or IoCS(7) are optionally reserved for the Centronics Port if used.
$\overline{\text{MemRdEnEven}}$	Output	Memory Read Enable for Even FCT245/543 Type Banks. On FCT260 type banks, it is the read enable for both even and odd banks.
$\overline{\text{MemRdEnOdd}}$	Output	Memory Read Enable for Odd FCT245/543 Type Banks. On FCT260 Type banks, it is the path select.
$\overline{\text{MemWrEnEven}}$	Output	Memory Write Enable for Even Banks.
$\overline{\text{MemWrEnOdd}}$	Output	Memory Write Enable for Odd Banks.
$\overline{\text{MemWrEn(3:0)}}$	Output	Memory Write Enable for each Byte Lane.
$\overline{\text{IoRdEn/DStrobe}}$	Output	I/O Read Enable or I/O Data Strobe.
$\overline{\text{IoWrEn/RdWr}}$	Output	I/O Write Enable or I/O Read/Write.

Table 1.1 R36100 Pin Descriptions

Pin Function and Description		
DMA Controller Pins		
$\overline{\text{DmaBusGnt}}(1:0)$	Output	DMA Bus Grant. Indicates that the CPU has tri-stated the bus and other DMA related signals.
$\overline{\text{DmaBusReq}}(1:0)$	Input	DMA Bus Request. Indicates that external DMA agent wants bus control.
$\overline{\text{DmaDone}}$	Input	DMA Transaction Done.
Serial Port Pins		
$\overline{\text{SerialPclkIn}}(1:0)$	Input	Optional Primary Serial Clock Input.
$\overline{\text{SerialSclk}}(1:0)$	I/O	Optional Secondary Serial Clock Input or Output.
$\overline{\text{SerialRxData}}(1:0)$	Input	Serial Receiver Data Stream.
$\overline{\text{SerialTxData}}(1:0)$	Output	Serial Transmitter Data Stream.
$\overline{\text{SerialCTS}}(1:0)$	Input	Serial Clear To Send.
$\overline{\text{SerialRTS}}(1:0)$	Output	Serial Request To Send.
$\overline{\text{SerialSync}}(1:0)$	I/O	Serial Frame Sync.
$\overline{\text{SerialDCD}}(1:0)$	Input	Serial Data Carrier Detect.
$\overline{\text{SerialDTR}}(1:0)$	Output	Serial Data Terminal Ready.
Timer Pins		
$\overline{\text{Timer TC}}(2:0)/$ $\overline{\text{TimerGate}}(2:0)$	I/O	Timer Terminal Count output or Timer Count Gate Enable input. Terminal Count asserts when Timer Count equals 0. Timer Gate enables Counter to count upward or to stop.
PIO Pins		
$\text{PIO}(31:0)$	I/O	Parallel Inputs or Parallel Outputs. Parallel inputs and parallel outputs are multiplexed with various peripheral inputs and peripheral outputs. If the peripheral is unused, the input or output pin can be reconfigured to be a general purpose input or output, respectively.
Bi-directional Centronics Interface Pins		
$\overline{\text{CentStrobe}}$	Input	Centronics Strobe. In Compatible mode, strobes data into the printer. Has other uses for other modes.
$\overline{\text{CentAck}}$	Output	Centronics Acknowledge. In Compatible mode, acknowledges a strobe. Has other uses for other modes.
$\overline{\text{CentBusy}}$	Output	Centronics Busy. In Compatible mode, delays the Host from sending more data. Has other uses for other modes.
$\overline{\text{CentPaperError}}$	Output	Centronics Paper Out/Jam Error. In Compatible mode, indicates that the printer has a paper error when asserted with $\overline{\text{CentFault}}$. Has other uses for other modes.
$\overline{\text{CentSelect}}$	Output	Centronics Select. In Compatible mode, used to indicate that this printer is on-line. Has other uses for other modes.
$\overline{\text{CentAutoFeed}}$	Input	Centronics Auto Page Feed. In Compatible mode, sends a paper feed to the printer. Has other uses for other modes.
$\overline{\text{CentInit}}$	Input	Centronics Initialization/Reset. In Compatible mode, resets the printer. Has other uses for other modes.

Table 1.1 R36100 Pin Descriptions (Continued)

Pin Function and Description		
Bi-directional Centronics Interface Pins (continued)		
$\overline{\text{CentFault}}$	Output	Centronics Fault. In Compatible mode, indicates that the printer has a problem. Has other uses for other modes.
$\overline{\text{CentSelectIn}}$	Input	Centronics Select In. In Compatible mode, indicates that the Host wants to select this printer on a shared cable. Has other uses for other modes.
CentHostStrobe	Output	Centronics Host Strobe. Used to latch Host data on the external FCT952/374 data transceiver during a Host write.
$\overline{\text{CentHostOEn}}$	Output	Centronics Host Output Enable. Used to enable the external FCT952/374 data transceiver during a Host read.
Laser Engine Interface Pins		
LaserVideoData	Output	Laser Video Data Stream.
LaserVideoClkIn	Input	Laser Video Clock Input. Accepts either the (1x) Video Data Stream frequency or 8 times (8x) the PLL frequency.
$\overline{\text{LaserLineSync}}$	Input	Laser Line Sync. Indicates that the laser drum is ready to start accepting data for a new line.
$\overline{\text{LaserPageSync}}$	Input	Laser Page Sync. Indicates that the laser drum is ready to start a new page.
Debug/Emulator Interface Pins		
JtagClkIn	Input	JTAG Clock Input (TCK). Test mode serial boundary scan input clock.
JtagModeSelect	Input	JTAG Mode Select (TSEL). Test mode serial boundary scan command data. In normal operating mode, $\overline{\text{JtagModeSelect}}$ should be left unasserted high.
JtagDataIn	Input	JTAG Data In (TDI). Test mode serial boundary scan register data input.
JtagDataOut	Output	JTAG Data Out (TDO). Test mode serial boundary scan register data output.
$\overline{\text{JtagReset}}$	Input	JTAG Reset (TRES*). Resets the JTAG test circuitry. Does not reset any other chip functions. In normal operating mode, $\overline{\text{JtagReset}}$ should be left asserted low.
Diagnostic Pins		
$\overline{\text{DiagC/UnC}}$	Output	Diagnostic Cached versus Uncached. On read bus transactions indicates whether the read is cached or uncached.
$\overline{\text{DiagInst/Data}}$	Output	Diagnostic Instruction versus Data. On read bus transactions indicates whether the read is for instructions or data.
$\overline{\text{DiagRun}}$	Output	Diagnostic Run. Indicates an internal pipeline run cycle. This pin has "pseudo-synchronous" timing.
$\overline{\text{DiagBranchTaken}}$	Output	Diagnostic Branch Taken. Indicates that a branch, jump, or exception has been taken. This pin has "pseudo-synchronous" timing.
$\overline{\text{DiagJRorExe}}$	Output	Diagnostic Jump Register or Exception occurring. Indicates that a jump register or exception is executing. This pin has "pseudo-synchronous" timing.
$\overline{\text{DiagInternalWr}}$	Output	Diagnostic Internal Write. Indicates that a MTCO to CP0 register \$3 is occurring.

Table 1.1 R36100 Pin Descriptions (Continued)

Pin Function and Description		
Diagnostic Pins (continued)		
$\overline{\text{DiagInstCacheWrDis}}$	Input	Diagnostic Cache Write Disable. Disables writes to the instruction and data cache. This pin has "pseudo-synchronous" timing and is not recommended for functional use.
$\overline{\text{DiagTriState}}$	Input	Diagnostic Tri-State all outputs. All outputs are tri-stated including $\overline{\text{SysClk}}$. This pin is asynchronous such that tri-stating asserts or de-asserts external output enables immediately.
$\overline{\text{DiagFCM}}$	Input	Diagnostic Force Cache Miss. This pin has "pseudo-synchronous" timing. If used for functional board tests, it is recommended that it be (de-)asserted statically at reset time and left (de-)asserted.
$\overline{\text{DiagIntDis}}$	Input	Diagnostic Interrupt Disable.
$\overline{\text{DiagNoCS}}$	Output	Diagnostic No Chip Select. No internal or external chip select has occurred for the current bus transaction, therefore an external state machine should handle the bus transaction.
$\overline{\text{DiagInternalDMA}}$	Output	Diagnostic Internal DMA. Asserts whenever any of the Internal DMA channels is generating the current bus transaction.
Exception Handling		
$\overline{\text{ExcSInt}}(2:0)$	Input	Exception Synchronized Interrupts. Also used as the reset initialization vector for 2:Boot16, 1:Boot8, and 0:BigEndian modes.
$\overline{\text{ExcInt}}(4:3)$	Input	Exception Interrupts.
$\overline{\text{ExcSBrCond}}(3:2)$	Input	Exception Synchronized Branch Condition inputs.
Power/Ground Pins		
VCC	Input	Power pin. All power pins must be connected. 5V or 3.3V depending on part type.
Gnd	Input	Ground pin (VSS). All ground pins must be connected. 0V.

Table 1.1 R36100 Pin Descriptions (Continued)

System Usage

The IDT R36100 RISController is specifically designed to easily implement low-cost memory systems. Typical low-cost memory systems use EPROMs and DRAM as well as application specific peripherals. Some embedded systems also optionally contain or substitute DRAM with static RAMs.

Figure 1.3 illustrates the low-system cost inherent in the R36100. For this example, which is typical of a low-cost laser printer, a 32-bit PROM interface is used, due to the size of the PDL interpreter. Other embedded systems could optionally use an 8-bit or 16-bit PROM, or an interleaved 64-bit interface.

A 16-bit font cartridge interface is provided through PCMCIA for add-in cards, and a 32-bit page buffer DRAM is used for high-resolution. In this example, a field or manufacturing upgrade to a larger page buffer is supported by the boot software and DRAM controller. Such a system features a very low entry price, with a range of field upgrade options. Note that the performance of the R36100 allows software frame buffer compression to be effective in reducing system DRAM while maintaining expected performance.

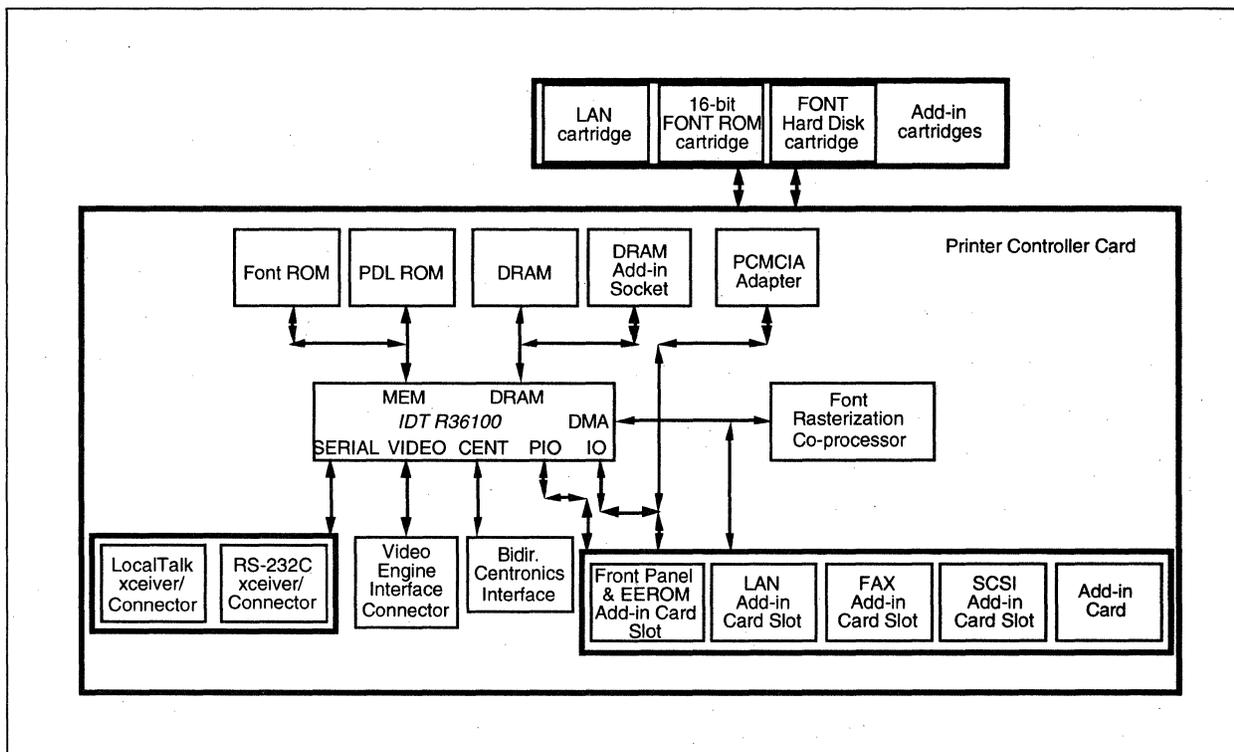


Figure 1.3 R36100-based Printer System

Development Support

The IDT R36100 RISController is supported by a rich set of development tools through the AdvantageIDT development tools program.

Figure 1.4 shows an overview of the system development process that is typically used with the R36100. Tools that allow timely, parallel development of hardware and software for R36100 family- based applications support all phases of R36100 project development.

These are some of the available support tools:

- Optimizing compilers from a number of leading compiler vendors.
- The IDT/c compiler, based on the GCC/GNU tool chain.
- The high-performance IDT floating point library software.
- The IDT Evaluation Board, which includes RAM, EPROM, I/O, and the IDT PROM Monitor.
- Adobe PostScript™ Page Description Language running on the IDT RISController family.
- The IDT/sim PROM Monitor, which implements a full PROM monitor (diagnostics, remote debug support, downloading utilities).
- IDT/kit (Kernel Integration Toolkit), providing library support and a frame work for the system run time environment.

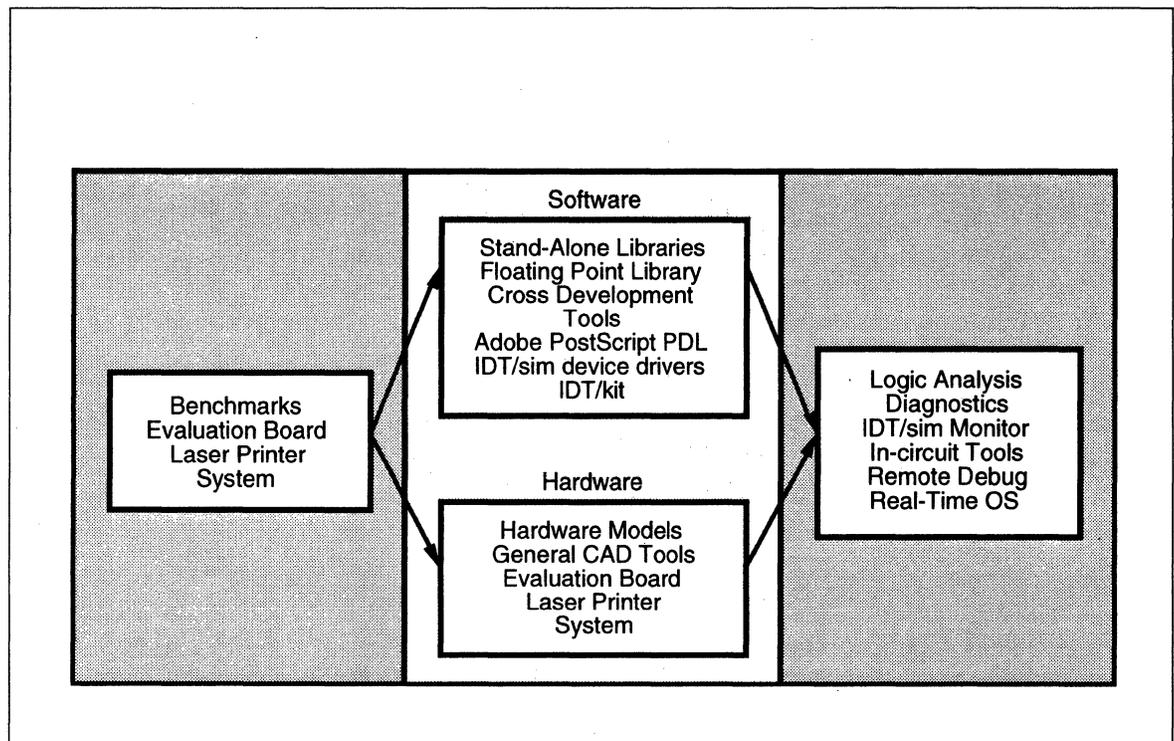
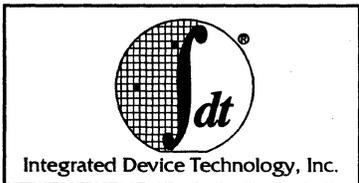


Figure 1.4 Development Support

Performance Overview

The R36100 RISController achieves a very high performance level that is based on:

- An efficient execution engine. The CPU executes almost all instructions at a single-cycle rate. Thus, the R36100 achieves over 24 dhrystone MIPS performance at 25MHz. By using a traditional 5-stage pipeline, the performance of the R36100 does not degrade in applications with a high-degree of data dependency.
- Large on-chip caches. The R36100 contains caches which are substantially larger than those on the majority of low-cost embedded microprocessors. These large caches minimize the number of bus transactions required and allow the R36100 to achieve actual sustained performance that is very close to its peak execution rate, even with low cost memory systems.
- Autonomous multiply and divide operations. The R36100 features an on-chip integer multiplier/divide unit which is separate from the other ALU. This allows the R36100 to perform multiply or divide operations in parallel with other integer operations, using a single multiply or divide instruction rather than with "step" operations.
- Integrated write buffer. The R36100 features a four deep write buffer, which captures store target addresses and data at the processor execution rate and retires it to main memory at the slower main memory access rate. Use of on-chip write buffers eliminates the need for the processor to stall when performing store operations.
- Burst read support. The R36100 enables the system designer to utilize page, static or nibble mode RAMs when performing read operations to minimize the main memory read penalty and increase the effective cache hit rates.
- Tightly coupled memory system. Integration of on-chip memory controllers allow system resources to be accessed and managed efficiently for the needs of the execution core.



Introduction

The IDT R36100 contains the same basic execution core as the IDT MIPS R3000A and the IDT R30xx RISControllers. In addition to being able to run software written for any of these processors, this enables the R36100 to achieve dramatic levels of performance, based on the efficiency of the execution engine.

This chapter gives an overview of the MIPS-I architecture implemented in the R36100, and discusses the programmers' model for this device. Further detail on the processor software model is found in the "IDT R30xx Family Software Reference Manual", available from IDT. The R36100 is fully ISA compatible with the R30xx family.

The R36100 is also address map compatible with the base versions of the R30xx family. However, to reduce system cost, the TLB functions present in the "E" versions are not available in the R36100; instead, the R36100 features were selected for minimal device and system cost.

Processor Features Overview

The R36100 has many of the same attributes of the IDT R30xx family, at a higher level of integration geared to lower system cost. These features include:

- **Full 32-bit Operation.** The R36100 contains thirty-two 32-bit general-purpose registers, and all instructions and addresses are 32 bits.
- **Efficient Pipelining.** The CPU utilizes a 5-stage pipeline design to achieve an execution rate approaching one instruction per cycle. Pipeline stalls, hazards, and exceptional events are handled precisely and efficiently.
- **Large On-Chip Instruction and Data Caches.** The R36100 utilizes large on-chip caches to provide high-bandwidth to the execution engine. The large size of the caches insures high hit rates, minimizing stalls due to cache miss processing and dramatically contributing to overall performance. Both the instruction and data cache can be accessed during a single CPU cycle.
- **On-chip Memory Management.** The R36100 is compatible with the base versions of the IDT R30xx family, which do not utilize a TLB, but perform fixed segment-based mapping of the virtual space to physical addresses. In addition, the R36100 allows kernel software to manage the system interface, by programming of the on-chip memory controllers and peripherals.

CPU Registers Overview

The IDT R36100 provides 32 general purpose 32-bit registers, an internal 32-bit Program Counter, and two dedicated 32-bit registers which hold the result of an integer multiply or divide operation. The CPU registers, illustrated in Figure 2.1, are discussed later in this chapter.

Note that the MIPS architecture does not use a traditional Program Status Word (PSW) register. The functions normally provided by such a register are instead provided through the use of “Set” instructions and conditional branches. By avoiding the use of traditional condition codes, the architecture can be more finely pipelined. This, coupled with the fine granularity of the instruction set, allows the compilers to achieve dramatically higher levels of optimizations than for traditional architectures.

Overflow and exceptional conditions are then handled through the use of the on-chip *Status* and *Cause* registers, which reside on-chip as part of the System Control Coprocessor (Coprocessor 0). These registers contain information about the run-time state of the machine, and any exceptional conditions it has encountered.

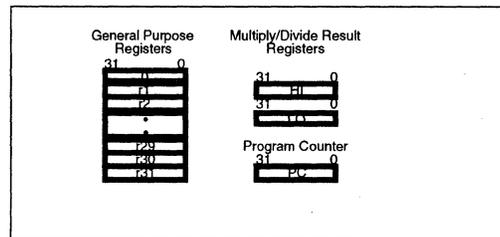


Figure 2.1 CPU Registers

Instruction Set Overview

All R36100 instructions are 32-bits long, and there are only three basic instruction formats. This approach dramatically simplifies instruction decoding, permitting higher frequency operation. More complicated (but less frequently used) operations and addressing modes are synthesized by the compiler/assembler, using sequences of the basic instruction set. This approach enables object code optimizations at a finer level of resolution than achievable in micro-coded CPU architectures.

Figure 2.2 shows the instruction set encoding used by the MIPS architecture. This approach simplifies instruction decoding in the CPU.

The R3000A instruction set (implemented in the R36100) can be divided into the following basic groups:

- **Load/Store** instructions move data between memory and the general registers. They are all encoded as “I-Type” instructions, and the only addressing mode implemented is base register plus signed, immediate offset. This directly enables the use of three distinct addressing modes: register plus offset; register direct; and immediate.
- **Computational** instructions perform arithmetic, logical, and shift operations on values in registers. They are encoded as either “R-Type” instructions, when both source operands as well as the result are general registers, and “I-Type”, when one of the source operands is a 16-bit immediate value. Computational instructions use a three address format, so that operations don’t needlessly interfere with the contents of source registers.
- **Jump and Branch** instructions change the control flow of a program. A Jump instruction can be encoded as a “J-Type” instruction, in which case the Jump target address is a paged absolute address formed by combining the 26-bit immediate value with the upper four bits of the Program Counter. This form is used for subroutine calls.

Alternately, Jumps can be encoded using the "R-Type" format, in which case the target address is a 32-bit value contained in one of the general registers. This form is typically used for returns and dispatches.

Branch operations are encoded as "I-Type" instructions. The target address is formed from a 16-bit displacement relative to the Program Counter.

The Jump and Link instructions save a return address in General Register r31. These are typically used as subroutine calls, where the subroutine return address is stored into r31 during the call operation.

- **Coprocessor** instructions perform operations on the co-processor set. Coprocessor Loads and Stores are always encoded as "I-Type" instructions; in the MIPS architecture, co-processor operational instructions have co-processor dependent formats.

In the R36100, the System Control Coprocessor (CPO) contains registers which are used in system interface control, cache control, and exception handling.

- **Special** instructions perform a variety of tasks, including movement of data between special and general registers, system calls, and breakpoint operations. They are always encoded as "R-Type" instructions.

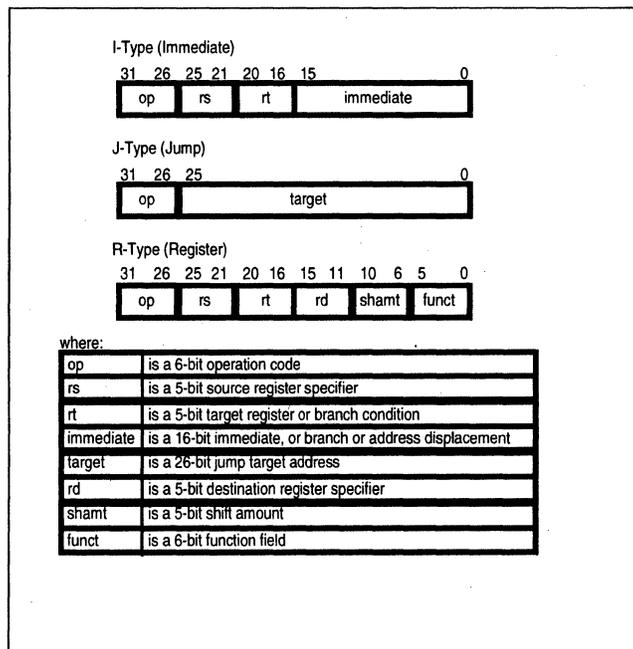


Figure 2.2 Instruction Encoding

OP	Description	OP	Description
	Load/Store Instructions		Multiply/Divide Instructions
LB	Load Byte	MULT	Multiply
LBU	Load Byte Unsigned	MULTU	Multiply Unsigned
LH	Load Halfword	DIV	Divide
LHU	Load Halfword Unsigned	DIVU	Divide Unsigned
LW	Load Word		
LWL	Load Word Left	MFHI	Move From HI
LWR	Load Word Right	MTHI	Move To HI
SB	Store Byte	MFLO	Move From LO
SH	Store Halfword	MTLO	Move To LO
SW	Store Word		
SWL	Store Word Left		Jump and Branch Instructions
SWR	Store Word Right	J	Jump
		JAL	Jump and Link
	Arithmetic Instructions (ALU Immediate)	JR	Jump to Register
		JALR	Jump and Link Register
ADDI	Add Immediate	BEQ	Branch on Equal
ADDIU	Add Immediate Unsigned	BNE	Branch on Not Equal
SLTI	Set on Less Than Immediate	BLEZ	Branch on Less than or Equal to Zero
SLTIU	Set on Less Than Immediate Unsigned	BGTZ	Branch on Greater Than Zero
ANDI	AND Immediate	BLTZ	Branch on Less Than Zero
ORI	OR Immediate	BGEZ	Branch on Greater Than or Equal to Zero
XORI	Exclusive OR Immediate	BLTZAL	Branch on Less Than Zero and Link
LUI	Load Upper Immediate	BGEZAL	Branch on Greater Than or Equal to Zero and Link
	Arithmetic Instructions (3-operand, register-type)		
ADD	Add		Special Instructions
ADDU	Add Unsigned	SYSCALL	System Call
SUB	Subtract	BREAK	Break
SUBU	Subtract Unsigned		
SLT	Set on Less Than		Coprocessor Instructions
SLTU	Set on Less Than Unsigned	LWCz	Load Word from Coprocessor
AND	AND	SWCz	Store Word to Coprocessor
OR	OR	MTCz	Move To Coprocessor
XOR	Exclusive OR	MFCz	Move From Coprocessor
NOR	NOR	CTCz	Move Control To Coprocessor
		CFCz	Move Control From Coprocessor
	Shift Instructions	COPz	Coprocessor Operation
SLL	Shift Left Logical	BCzT	Branch on Coprocessor z True
SRL	Shift Right Logical	BCzF	Branch on Coprocessor z False
SRA	Shift Right Arithmetic		
SLLV	Shift Left Logical Variable		System Control Coprocessor (CPO) Instructions
SRLV	Shift Right Logical Variable		
SRAV	Shift Right Arithmetic Variable	MTC0	Move To CPO
		MFC0	Move From CPO
		TLBR†	Read indexed TLB entry
		TLBWI†	Write indexed TLB entry
		TLBWR†	Write Random TLB entry
		TLBP†	Probe TLB for matching entry
		RFE	Restore From Exception

†These instructions are not valid with the R36100, which does not include a TLB.

Table 2.1 Instruction Set Mnemonics

Table 2.1 lists the instruction set mnemonics of the R36100. More detail on these operations is presented later in this chapter. For further detail, consult the "IDT R30xx Family Software Reference Manual", available from IDT.

Programming Model

This section describes the organization of data in the general registers and in memory, and discusses the set of general registers available. A summary description of all of the CPU registers is presented.

Data Formats and Addressing

The MIPS-I architecture defines a word as 32-bits, a half-word as 16-bits, and a byte as 8-bits. The byte ordering convention is configurable during hardware reset into either a *big-endian* or *little-endian* convention.

When configured as a big-endian system, byte 0 is always the most significant (leftmost) byte in a word. But when configured as a little-endian system, byte 0 is always the least significant (rightmost) byte in a word.

Figure 2.3 shows the ordering of bytes within words and the ordering of words within multiple word structures for the big-endian and little-endian conventions.

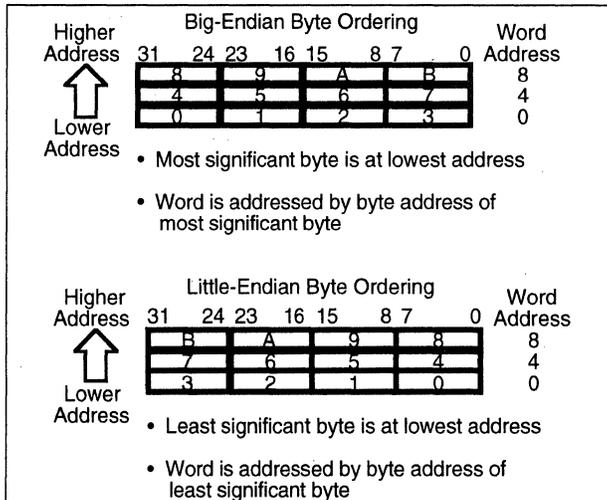


Figure 2.3 Byte Ordering Conventions

The R36100 uses byte addressing for all accesses, including half-word and word. The MIPS architecture has alignment constraints that require half-word accesses to be aligned on an even byte boundary, and word accesses to be aligned on a modulo-4 byte boundary. Thus, in big-endian systems, the address of a multiple-byte data item is the address of the most-significant byte, while in little-endian systems it is the address of the least-significant byte of the structure.

The MIPS instruction set provides special instructions for addressing 32-bit words which are not aligned on 4-byte boundaries. These instructions, which are Load/Store Left/Right, are used in pairs to provide addressing of misaligned words. This effectively means that these types of data movements require only one-additional instruction cycle over that required for properly aligned words (note that unaligned data is read by the CPU in the same number of cycles as would be required for a full hardware solution, and provides a much more efficient way of dealing with this case than is possible using sequences of loads/stores and shift operations or by using traps. Various tool chains, such as the IDT/c compiler, can automatically generate these instructions for "packed" data. Figure 2.4 shows the bytes accessed when addressing a mis-aligned word with a byte address of 3, for each of the two byte ordering conventions.

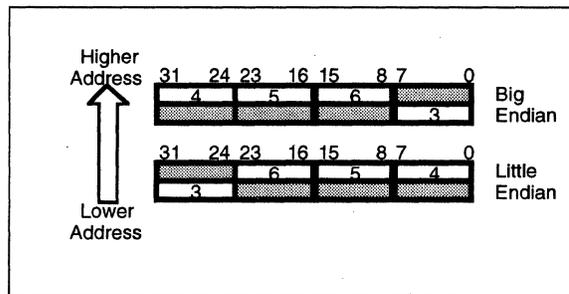


Figure 2.4 Unaligned Words

CPU General Registers

The R36100 contains 32 general registers, each containing a single 32-bit word. The 32 general registers are treated symmetrically (orthogonally), with two notable exceptions: general register r0 is hardwired to a zero value, and r31 is used as the link register in Jump and Link instructions.

Register r0 maintains the value zero under all conditions, when used as a source register, and discards data written to it. Thus, instructions which attempt to write to it may be used as No-Op Instructions. The use of a register wired to the zero value allows the simple synthesis of different addressing modes, no-ops, register or memory clear operations, etc., without requiring expansion of the basic instruction set.

Register r31 is used as the link register in jump and link instructions. These instructions are used in subroutine calls, and the subroutine return address is placed in register r31. This register can be written to or read as a normal register in other operations.

In addition to the general registers, the CPU contains two registers (HI and LO) which store the double-word, 64-bit result of integer multiply operations, and the quotient and remainder of integer divide operations.

CP0 Special Registers

In addition to the general CPU registers, the R36100 contains a number of special registers on-chip. These registers logically reside in the on-chip System Control Co-processor CP0, and are used in memory management and exception handling.

Table 2.2 on page 8 shows the logical CP0 address of each of the registers. The format of each of these registers, and their use, is discussed in later chapters. Note that the MIPS architecture allows CP0 to vary by implementation; the R36100 contains some new CP0 registers not found in other R30xx family members; however, their definition is such that it still remains possible to use a single binary program across all family members, in that these registers are typically managed only at reset.

Number	Mnemonic	Description
0	Reserved(1)	
1	Reserved(1)	
2	Reserved(1)	
3	Config(3)	Cache Usage Configuration
4	Reserved(1)	
5-7	Reserved	
8	BadVAddr	Bad Virtual Address
9	Reserved(3)	
10	Reserved(3)	
11	Reserved(3)	
12	SR	Status Register
13	Cause	Cause of Last Exception
14	EPC	Exception Program Counter
15	PRId	Processor Revision Identifier
16-31	Reserved	
<p>NOTES:</p> <ol style="list-style-type: none"> 1. This register is used in Extended Architecture CPUs to control the TLB and virtual memory system. In the "E" versions, register \$2 is "TLB EntryLo", and register \$10 is "TLB EntryHi". 2. This register is reserved in other family members. 3. This register has a different meaning in other family members. 		

Table 2.2 R36100 CP0 Registers

Operating Modes

The R36100 supports two different operating modes: *User* and *Kernel* modes. The R36100 normally operates in User mode until an exception is detected, forcing it into kernel mode. It remains in Kernel mode until a Return From Exception (RFE) instruction is executed, returning it to its previous operation mode.

The processor supports these levels of protection by segmenting the 4GB virtual address space into 4 distinct segments. One segment is accessible from either the User state or the Kernel mode, and the other three segments are only accessible from kernel mode.

In addition to providing memory address protection, the kernel can protect the co-processors (in the case of the R36100, CP0) from access or modification by the user task. Chapter 4 discusses the memory management facilities of the processor.

Pipeline Architecture

The IDT R36100 uses the same basic pipeline structure as that implemented in the R3000A. Thus, the execution of a single instruction is performed in the following five distinct stages:

- Instruction Fetch (IF). In this stage, the instruction virtual address is translated to a physical address and the instruction is read from the internal Instruction Cache.
- Read (RD). During this stage, the instruction is decoded and required operands are read from the on-chip register file.
- ALU. The required operation is performed on the instruction operands.
- Memory Access (MEM). If the instruction was a load or store, the Data Cache is accessed. Note that there is a skew between the instruction cycle which fetches the instruction and the one in which the required data transfer occurs. This skew is a result of the intervening pipe-stages.
- Write Back (WB). During the write back pipestage, the results of the ALU stage operation are updated into the on-chip register file.

Each of these pipestages require approximately one CPU cycle, as shown in Figure 2.5. Parts of some operations lap into the next cycle, while other operations require only 1/2 cycle.

The net effect of the pipeline structure is that a new instruction can be initiated every clock cycle. Thus, the execution of five instructions at a time is overlapped, as shown in Figure 2.6.

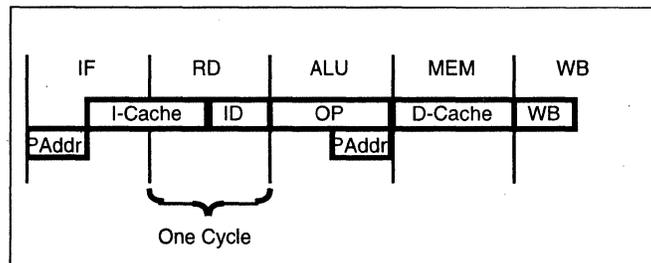


Figure 2.5 5-Stage Pipeline

The pipeline operates efficiently, because different CPU resources such as address and data bus access, ALU operations, and the register file, are utilized on a non-interfering basis.

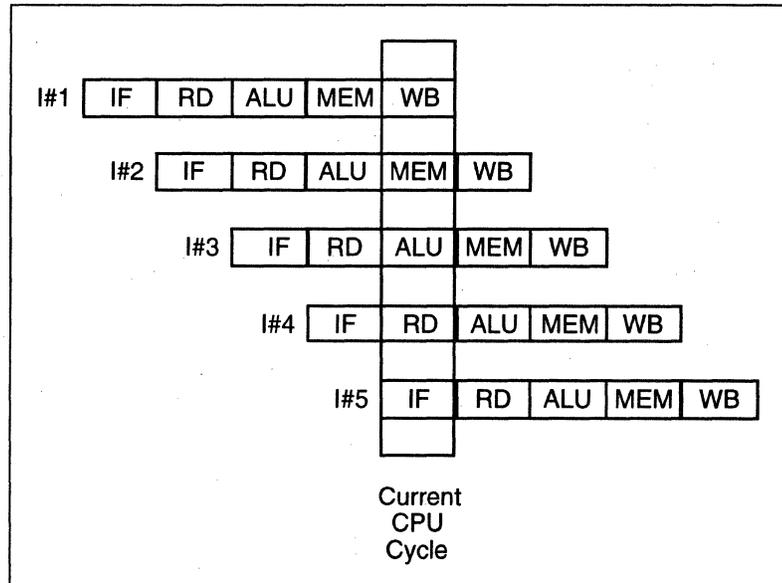


Figure 2.6 5-Instructions per Clock Cycle

Pipeline Hazards

In a pipelined machine such as the R36100, there are certain instructions which, based on the pipeline structure, can potentially disrupt the smooth operation of the pipeline. The basic problem is that the current pipestage of an instruction may require the result of a previous instruction, still in the pipeline, whose result is not yet available. This class of problems is referred to as pipeline hazards.

An example of a potential pipeline hazard occurs when a computational instruction ($n+1$) requires the result of the immediately prior instruction (instruction n). Instruction $n+1$ wants to access the register file during the RF pipestage. However, instruction n has not yet completed its register writeback operation, and thus the current value is not available directly from the register file. In this case, special logic within the execution engine forwards the result of instruction n 's ALU operation to instruction $n+1$, prior to the true writeback operation. The pipeline is undisturbed, and no pipeline *stalls* need to occur.

Another example of a pipeline hazard handled in hardware is the integer multiply and divide operations. If an instruction attempts to access the HI or LO registers prior to the completion of the multiply or divide, that instruction will be *interlocked* (held off) until the multiply or divide operation completes. Thus, the programmer is isolated from the actual execution time of this operation. The optimizing compilers attempt to schedule as many instructions as possible between the start of the multiply/divide and the access of its result, to minimize stalls.

However, not all pipeline hazards are handled in hardware. There are two notable categories of instructions which require software intervention to insure logical operation. The optimizing compilers (and peephole scheduler of the assembler) are capable of insuring proper execution. These two instruction classes are:

- Load instructions have a delay, or latency, of one cycle before the data loaded from memory is available another instruction. This is because the ALU stage of the immediately subsequent instruction is processed simultaneously with the Data Cache access of the load operation. Figure 2.7 illustrates the cause of this delay slot.

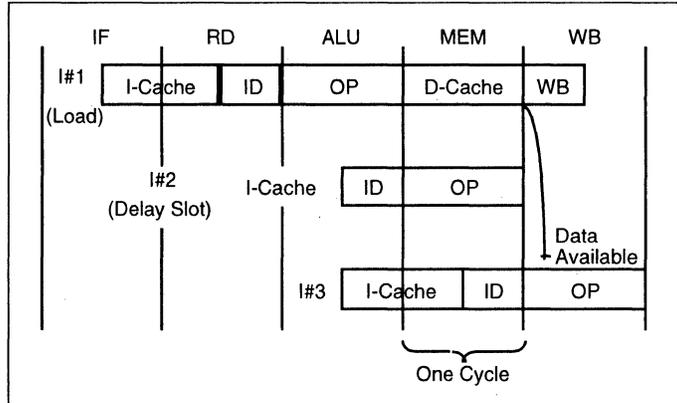


Figure 2.7 Load Delay

- Jump and Branch instructions have a delay of one cycle before the program flow change can occur. This is due to the fact that the next instruction is fetched prior to the decode and ALU stage of the jump/branch operation. Figure 2.8 illustrates the cause of this delay slot.

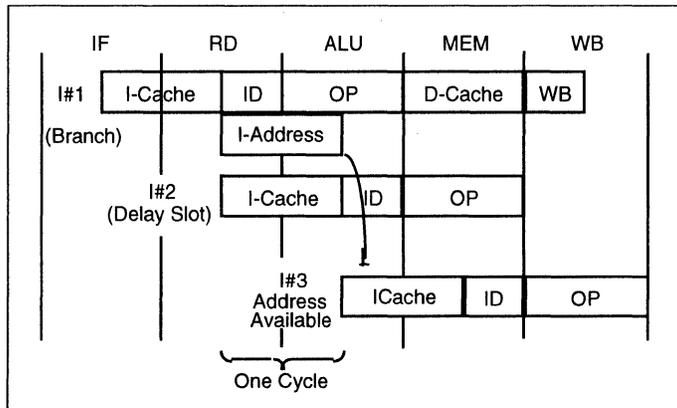


Figure 2.8 Branch Delay

The R36100 continues execution, despite the delay in the operation. Thus, loads, jumps and branches do not disrupt the pipeline flow of instructions, and the processor always executes the instruction immediately following one of these “delayed” instructions.

Note: Note that there may also be latencies associated with changes to various of the CPO registers; for example, changing the bus interface control register may require multiple cycles before the change is actually reflected in the chip interface.

Rather than include extensive pipeline control logic, the MIPS-I instruction set gives responsibility for dealing with “delay slots” to software. Thus, peephole optimizations (which can be performed as part of compilation or assembly) can re-order the code to insure that the instruction in the delay slot does not require the logical result of the “delayed” instruction. In the worst case, a NOP can be inserted to guarantee proper software execution.

Chapter 6 discusses the impact of pipelining on exception handling. In general, when an instruction causes an exception, it is desirable for all instructions initiated prior to that instruction to complete, and all subsequent instructions to abort. This insures that the machine state presented to the exception handler reflects the logical state that existed at the time the exception was detected. In addition, it is desirable to avoid requiring software to explicitly manage the pipeline when handling or returning from exceptions. The IDT R36100 pipeline is designed to properly manage exceptional events.

Instruction Set Summary

This section provides an overview of the R36100 instruction set by presenting each category of instructions in a tabular summary form. Refer to the “IDT R30xx Family Software Reference Manual”, for a detailed description of each instruction.

Instruction Formats

Every instruction consists of a single word (32 bits) aligned on a word boundary. There are only three instruction formats, as shown in Figure 2.2 on page 3. This approach simplifies instruction decoding. More complicated or less frequently used operations and addressing modes are synthesized by the compilers.

Instruction Notational Conventions

In this manual, all variable sub-fields in an instruction format (such as *rs*, *rt*, *immediate*, and so on) are shown in lower-case names.

For the sake of clarity, an alias is sometimes used for a variable sub-field in the formats of specific instructions. For example, “*base*” rather than “*rs*” is used in the format for Load and Store instructions. Such an alias is always lower case, since it refers to a variable sub-field.

Instruction opcodes are shown in upper case.

The actual bit encoding for all the mnemonics is specified at the end of this chapter.

Load and Store Instructions

Load/Store instructions move data between memory and general registers. They are all I-type instructions. The only addressing mode directly supported is base register plus 16-bit signed immediate offset. This can be used to directly implement immediate addressing (using the r0 register) or register direct (using an immediate offset value of zero).

All load operations have a latency effect of one instruction. That is, the data being loaded from memory into a register is not available to the instruction that immediately follows the load instruction: the data is available to the second instruction after the load instruction. An exception to this rule is that for the target register for the “load word left” and “load word right” instructions may be specified as the same register used as the destination of the related unaligned load instruction that immediately precedes it.

The Load/Store instruction opcode determines the size of the data item to be loaded or stored, as shown in Table 2.1 on page 4. Regardless of access type or byte numbering-order (endian-ness), the address specifies the byte which has the smallest byte address of all bytes in the addressed field. For a big-endian access, this is the most significant byte; for a little-endian access, this is the least significant byte. Note that in an R36100 based system, the endianness of a given access is dynamic, in that the RE (Reverse Endianness) bit of the Status Register can be used to force user space accesses of the opposite byte convention of the kernel.

Big-Endian (32-bit memory system)

Size	CPU Core VAdrLo(1)	CPU Core VAdrLo(0)	BE(3) Data(31:24)	BE(2) Data(23:16)	BE(1) Data(15:8)	BE(0) Data(7:0)
Word	0	0	Yes	Yes	Yes	Yes
Tri-Byte	0	0	Yes	Yes	Yes	No
Tri-Byte	0	1	No	Yes	Yes	Yes
16-Bit	0	0	Yes	Yes	No	No
16-Bit	1	0	No	No	Yes	Yes
Byte	0	0	Yes	No	No	No
Byte	0	1	No	Yes	No	No
Byte	1	0	No	No	Yes	No
Byte	1	1	No	No	No	Yes

Table 2.3 Big-Endian (32-bit memory system)

Little-Endian (32-bit memory system)

Size	VAdrLo(1)	VAdrLo(0)	BE(3) Data(31:24)	BE(2) Data(23:16)	BE(1) Data(15:8)	BE(0) Data(7:0)
Word	0	0	Yes	Yes	Yes	Yes
Tri-Byte	0	0	No	Yes	Yes	Yes
Tri-Byte	0	1	Yes	Yes	Yes	No
16-Bit	0	0	No	No	Yes	Yes
16-Bit	1	0	Yes	Yes	No	No
Byte	0	0	No	No	No	Yes
Byte	0	1	No	No	Yes	No
Byte	1	0	No	Yes	No	No
Byte	1	1	Yes	No	No	No

Table 2.4 Byte Addressing in Load/Store Operations (32-bit memory)

Big-Endian (16-bit memory system)

Size	CPU Core VAdrLo(1)	CPU Core VAdrLo(0)	First Transfer		Second Transfer	
			BE16(1) Data(31:24)	BE16(0) Data(23:16)	BE16(1) Data(31:24)	BE16(0) Data(23:16)
Word	0	0	Yes	Yes	Yes	Yes
Tri-Byte	0	0	Yes	Yes	Yes	No
Tri-Byte	0	1	No	Yes	Yes	Yes
16-Bit	0	0	Yes	Yes	N/A	N/A
16-Bit	1	0	Yes	Yes	N/A	N/A
Byte	0	0	Yes	No	N/A	N/A
Byte	0	1	No	Yes	N/A	N/A
Byte	1	0	Yes	No	N/A	N/A
Byte	1	1	No	Yes	N/A	N/A

Table 2.5 Big-Endian (16-bit memory system)

Little-Endian (16-bit memory system)

Size	CPU Core VAdrLo(1)	CPU Core VAdrLo(0)	First Transfer		Second Transfer	
			BE16(1) Data(15:8)	BE16(0) Data(7:0)	BE16(1) Data(15:8)	BE16(0) Data(7:0)
Word	0	0	Yes	Yes	Yes	Yes
Tri-Byte	0	0	Yes	Yes	No	Yes
Tri-Byte	0	1	Yes	No	Yes	Yes
16-Bit	0	0	Yes	Yes	N/A	N/A
16-Bit	1	0	Yes	Yes	N/A	N/A
Byte	0	0	No	Yes	N/A	N/A
Byte	0	1	Yes	No	N/A	N/A
Byte	1	0	No	Yes	N/A	N/A
Byte	1	1	Yes	No	N/A	N/A

**Table 2.6 Byte Addressing in Load/Store Operations
(16-bit memory)**

Note that the size of the operand requested by the load instruction is independent of the memory width of the addressed memory. Thus, if the actual size of the datum is 32-bits, software can safely use a load or store word instruction, even if the addressed memory is actually only 8- or 16-bits wide. The bus interface unit will interact with CPO to determine the width of the addressed memory, and will, if necessary, perform multiple datum transfers to satisfy a single load or store instruction.

The bytes within the addressed word that are used can be determined directly from the access size and the two low-order bits of the address, as shown in Table 2.3, Table 2.4, Table 2.5, and Table 2.6. Note that certain combinations of access type and low-order address bits can never occur: only the combinations shown in these tables are permissible.

Table 2.7 shows the load/store instructions supported by the MIPS-I ISA.

Instruction	Format and Description
Load Byte	LB rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Sign-extend contents of addressed byte and load into rt.
Load Byte Unsigned	LBU rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Zero-extend contents of addressed byte and load into rt.
Load Halfword	LH rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Sign-extend contents of addressed half-word and load into rt.
Load Halfword Unsigned	LHU rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Zero-extend contents of addressed half-word and load into rt.
Load Word	LW rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Load contents of addressed word into register rt.
Load Word Left	LWL rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Shift addressed word left so that addressed byte is leftmost byte of a word. Merge bytes from memory with contents of register rt and load result into register rt.
Load Word Right	LWR rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Shift addressed word right so that addressed byte is rightmost byte of a word. Merge bytes from memory with contents of register rt and load result into register rt.
Store Byte	SB rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Store least significant byte of register rt at addressed location.
Store Halfword	SH rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Store least significant halfword of register rt at addressed location.
Store Word	SW rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Store least significant word of register rt at addressed location.
Store Word Left	SWL rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Shift contents of register rt right so that leftmost byte of the word is in position of addressed byte. Store bytes containing original data into corresponding bytes at addressed byte.
Store Word Right	SWR rt, offset (base) Sign-extend 16-bit offset and add to contents of register base to form address. Shift contents of register rt left so that rightmost byte of the word is in position of addressed byte. Store bytes containing original data into corresponding bytes at addressed byte.

Table 2.7 Load and Store Instructions

Computational Instructions

Computational instructions perform arithmetic, logical and shift operations on values in registers. They occur in both R-type (both operands are registers) and I-type (one operand is a 16-bit immediate) formats. There are four categories of computational instructions:

- **ALU Immediate** instructions are summarized in Table 2.8.
- **3-Operand Register-Type** instructions are summarized in Table 2.9 on page 18.
- **Shift** instructions are summarized in Table 2.10 on page 19.
- **Multiply/Divide** instructions are summarized in Table 2.11 on page 19.

Instruction	Format and Description
ADD Immediate	ADDI rt, rs, immediate Add 16-bit sign-extended immediate to register rs and place 32-bit result in register rt. Trap on two's complement overflow.
ADD Immediate Unsigned	ADDIU rt, rs, immediate Add 16-bit sign-extended immediate to register rs and place 32-bit result in register rt. Do not trap on overflow.
Set on Less Than Immediate	SLTI rt, rs, immediate Compare 16-bit sign-extended immediate with register rs as signed 32-bit integers. Result = 1 if rs is less than immediate; otherwise result = 0. Place result in register rt.
Set on Less Than Unsigned Immediate	SLTIU rt, rs, immediate Compare 16-bit sign-extended immediate with register rs as unsigned 32-bit integers. Result = 1 if rs is less than immediate; otherwise result = 0. Place result in register rt. Do not trap on overflow.
AND Immediate	ANDI rt, rs, immediate Zero-extend 16-bit immediate, AND with contents of register rs and place result in register rt.
OR Immediate	ORI rt, rs, immediate Zero-extend 16-bit immediate, OR with contents of register rs and place result in register rt.
Exclusive OR Immediate	XORI rt, rs, immediate Zero-extend 16-bit immediate, exclusive OR with contents of register rs and place result in register rt.
Load Upper Immediate	LUI rt, immediate Shift 16-bit immediate left 16 bits. Set least significant 16 bits of word to zeroes. Store result in register rt.

Table 2.8 ALU Immediate Operations

Instruction	Format and Description
Add	ADD rd, rs, rt Add contents of registers rs and rt and place 32-bit result in register rd. Trap on two's complement overflow.
ADD Unsigned	ADDU rd, rs, rt Add contents of registers rs and rt and place 32-bit result in register rd. Do not trap on overflow.
Subtract	SUB rd, rs, rt Subtract contents of registers rt and rs and place 32-bit result in register rd. Trap on two's complement overflow.
Subtract Unsigned	SUBU rd, rs, rt Subtract contents of registers rt and rs and place 32-bit result in register rd. Do not trap on overflow.
Set on Less Than	SLT rd, rs, rt Compare contents of register rt to register rs (as signed 32-bit integers). If register rs is less than rt, result = 1; otherwise, result = 0.
Set on Less Than Unsigned	SLTU rd, rs, rt Compare contents of register rt to register rs (as unsigned 32-bit integers). If register rs is less than rt, result = 1; otherwise, result = 0.
AND	AND rd, rs, rt Bit-wise AND contents of registers rs and rt and place result in register rd.
OR	OR rd, rs, rt Bit-wise OR contents of registers rs and rt and place result in register rd.
Exclusive OR	XOR rd, rs, rt Bit-wise Exclusive OR contents of registers rs and rt and place result in register rd.
NOR	NOR rd, rs, rt Bit-wise NOR contents of registers rs and rt and place result in register rd.

Table 2.9 Three Operand Register-Type Operations

Instruction	Format and Description
Shift Left Logical	SLL rd, rt, shamt Shift contents of register rt left by shamt bits, inserting zeroes into low order bits. Place 32-bit result in register rd.
Shift Right Logical	SRL rd, rt, shamt Shift contents of register rt right by shamt bits, inserting zeroes into high order bits. Place 32-bit result in register rd.
Shift Right Arithmetic	SRA rd, rt, shamt Shift contents of register rt right by shamt bits, sign-extending the high order bits. Place 32-bit result in register rd.
Shift Left Logical Variable	SLLV rd, rt, rs Shift contents of register rt left. Low-order 5 bits of register rs specify number of bits to shift. Insert zeroes into low order bits of rt and place 32-bit result in register rd.
Shift Right Logical Variable	SRLV rd, rt, rs Shift contents of register rt right. Low-order 5 bits of register rs specify number of bits to shift. Insert zeroes into high order bits of rt and place 32-bit result in register rd.
Shift Right Arithmetic Variable	SRAV rd, rt, rs Shift contents of register rt right. Low-order 5 bits of register rs specify number of bits to shift. Sign-extend the high order bits of rt and place 32-bit result in register rd.

Table 2.10 Shift Operations

Instruction	Format and Description
Multiply	MULT rs, rt Multiply contents of registers rs and rt as twos complement values. Place 64-bit result in special registers HI/LO
Multiply Unsigned	MULTU rs, rt Multiply contents of registers rs and rt as unsigned values. Place 64-bit result in special registers HI/LO
Divide	DIV rs, rt Divide contents of register rs by rt treating operands as twos complements values. Place 32-bit quotient in special register LO, and 32-bit remainder in HI.
Divide Unsigned	DIVU rs, rt Divide contents of register rs by rt treating operands as unsigned values. Place 32-bit quotient in special register LO, and 32-bit remainder in HI.
Move From HI	MFHI rd Move contents of special register HI to register rd.
Move From LO	MFLO rd Move contents of special register LO to register rd.
Move To HI	MTHI rd Move contents of special register rd to special register HI.
Move To LO	MTLO rd Move contents of register rd to special register LO.

Table 2.11 Multiply and Divide Operations

Jump and Branch instructions

Jump and Branch instructions change the control flow of a program. All Jump and Branch instructions occur with a one instruction delay: that is, the instruction immediately following the jump or branch is always executed while the target instruction is being fetched, regardless of whether the branch is to be taken.

An assembler has several possibilities for utilizing the branch delay slot productively:

- It can insert an instruction that logically precedes the branch instruction in the delay slot since the instruction immediately following the jump/branch effectively belongs to the block preceding the transfer instruction.
- It can replicate the instruction that is the target of the branch/jump into the delay slot provided that no side-effects occur if the branch falls through.
- It can move an instruction up from below the branch into the delay slot, provided that no side-effects occur if the branch is taken.
- If no other instruction is available, it can insert a NOP instruction in the delay slot.

The J-type instruction format is used for both jumps-and-links for subroutine calls. In this format, the 26-bit target address is shifted left two bits, and combined with high-order 4 bits of the current program counter to form a 32-bit absolute address.

The R-type instruction format which takes a 32-bit byte address contained in a register is used for returns, dispatches, and cross-page jumps.

Branches have 16-bit offsets relative to the program counter (I-type). Jump-and-Link and Branch-and-Link instructions save a return address in register r31.

Table 2.12 summarizes the R36100's Jump instructions and Table 2.13 on page 21 summarizes the Branch instructions.

Instruction	Format and Description
Jump	J target Shift 26-bit target address left two bits, combine with high-order 4 bits of PC and jump to address with a one instruction delay.
Jump and Link	JAL target Shift 26-bit target address left two bits, combine with high-order 4 bits of PC and jump to address with a one instruction delay. Place address of instruction following delay slot in r31 (link register).
Jump Register	JR rs Jump to address contained in register rs with a one instruction delay.
Jump and Link Register	JALR rs, rd Jump to address contained in register rs with a one instruction delay. Place address of instruction following delay slot in rd.

Table 2.12 Jump Instructions

Instruction	Format and Description
	Branch Target: All Branch instruction target addresses are computed as follows: Add address of instruction in delay slot and the 16-bit offset (shifted left two bits and sign-extended to 32 bits). All branches occur with a delay of one instruction.
Branch on Equal	BEQ rs, rt, offset Branch to target address if register rs equal to rt
Branch on Not Equal	BNE rs, rt, offset Branch to target address if register rs not equal to rt.
Branch on Less than or Equal Zero	BLEZ rs, offset Branch to target address if register rs less than or equal to 0.
Branch on Greater Than Zero	BGTZ rs, offset Branch to target address if register rs greater than 0.
Branch on Less Than Zero	BLTZ rs, offset Branch to target address if register rs less than 0.
Branch on Greater than or Equal Zero	BGEZ rs, offset Branch to target address if register rs greater than or equal to 0.
Branch on Less Than Zero And Link	BLTZAL rs, offset Place address of instruction following delay slot in register r31 (link register). Branch to target address if register rs less than 0.
Branch on greater than or Equal Zero And Link	BGEZAL rs, offset Place address of instruction following delay slot in register r31 (link register). Branch to target address if register rs is greater than or equal to 0.

Table 2.13 Branch Instructions

Special Instructions

The two Special instructions let software initiate traps. They are always R-type. Table 2.14 summarizes Special Instructions.

Instruction	Format and Description
System Call	SYSCALL Initiates system call trap, immediately transferring control to exception handler.
Breakpoint	BREAK Initiates breakpoint trap, immediately transferring control to exception handler.

Table 2.14 Special Instructions

Co-processor Instructions

Co-processor instructions perform operations in the co-processors. Co-processor Loads and Stores are I-type. Co-processor computational instructions have co-processor-dependent formats. The only co-processor operations of relevance for the R36100 are those targeted at the on-chip CPO.

Table 2.15 summarizes the Co-processor Instruction Set of the MIPS ISA.

Instruction	Format and Description
Load Word to Co-processor	LWCz rt, offset (base) Sign-extend 16-bit offset and add to base to form address. Load contents of addressed word into co-processor register rt of co-processor unit z.
Store Word from Co-processor	SWCz rt, offset (base) Sign-extend 16-bit offset and add to base to form address. Store contents of co-processor register rt from co-processor unit z at addressed memory word.
Move To Co-processor	MTCz rt, rd Move contents of CPU register rt into co-processor register rd of co-processor unit z.
Move from Co-processor	MFCz rt, rd Move contents of co-processor register rd from co-processor unit z to CPU register rt.
Move Control To Co-processor	CTCz rt, rd Move contents of CPU register rt into co-processor control register rd of co-processor unit z.
Move Control From Co-processor	CFCz rt, rd Move contents of control register rd of co-processor unit z into CPU register rt.
Co-processor Operation	COPz cofun Co-processor z performs an operation. The state of the R36100 is not modified by a co-processor operation.
Branch on Co-processor z True	BCzT offset Compute a branch target address by adding address of instruction in the 16-bit offset (shifted left two bits and sign-extended to 32-bits). Branch to the target address (with a delay of one instruction) if co-processor z's condition line is true.
Branch on Co-processor z False	BCzF offset Compute a branch target address by adding address of instruction in the 16-bit offset (shifted left two bits and sign-extended to 32-bits). Branch to the target address (with a delay of one instruction) if co-processor z's condition line is false.

Table 2.15 Co-Processor Operations

System Control Co-processor (CPO) Instructions

Co-processor 0 instructions perform operations on the System Control Co-processor (CPO) registers to manipulate the memory management, bus programmability, timer, and exception handling facilities of the processor. Memory management, bus programmability, and exception handling are described in later chapters.

Table 2.16 summarizes the instructions available to work with CPO.

Instruction	Format and Description
Move To CPO	MTCO rt, rd Store contents of CPU register rt into register rd of CPO. This follows the convention of store operations.
Move From CPO	MFCO rt, rd Load CPU register rt with contents of CPO register rd.
Read Indexed TLB Entry	TLBR† Load EntryHi and EntryLo registers with TLB entry pointed at by Index register.
Write Indexed TLB Entry	TLBWI† Load TLB entry pointed at by Index register with contents of EntryHi and EntryLo registers.
Write Random TLB Entry	TLBWR† Load TLB entry pointed at by Random register with contents of EntryHi and EntryLo registers.
Probe TLB for Matching Entry	TLBP† Load Index register with address of TLB entry whose contents match EntryHi and EntryLo. If no TLB entry matches, set high-order bit of Index register.
Restore From Exception	RFE Restore previous interrupt mask and mode bits of status register into current status bits. Restore old status bits into previous status bits.
†These operations are undefined/reserved in the R36100, which does not include an on-chip TLB.	

Table 2.16 System Control Co-Processor (CPO) Operations

R36100 Opcode Encoding

Table 2.17 shows the opcode encoding for the MIPS architecture.

		OPCODE							
		0	1	2	3	4	5	6	7
31..29	28..26								
0	0	SPECIAL	BCOND	J	JAL	BEQ	BNE	BLEZ	BGTZ
1	1	ADDI	ADDIU	SLTI	SLTIU	ANDI	ORI	XORI	LUI
2	2	COP0	COP1	COP2	COP3	†	†	†	†
3	3	†	†	†	†	†	†	†	†
4	4	LB	LH	LWL	LW	LBU	LHU	LWR	†
5	5	SB	SH	SWL	SW	†	†	SWR	†
6	6	LWC0	LWC1	LWC2	LWC3	†	†	†	†
7	7	SWC0	SWC1	SWC2	SWC3	†	†	†	†
		SPECIAL							
		0	1	2	3	4	5	6	7
5..3	2..0								
0	0	SLL	†	SRL	SRA	SLLV	†	SRLV	SRAV
1	1	JR	JALR	†	†	SYSCALL	BREAK	†	†
2	2	MFHI	MTHI	MFLO	MTLO	†	†	†	†
3	3	MULT	MULTU	DIV	DIVU	†	†	†	†
4	4	ADD	ADDU	SUB	SUBU	AND	OR	XOR	NOR
5	5	†	†	SLT	SLTU	†	†	†	†
6	6	†	†	†	†	†	†	†	†
7	7	†	†	†	†	†	†	†	†
		BCOND							
		0	1	2	3	4	5	6	7
20..19	18..16								
0	0	BLTZ	BGEZ						
1	1								
2	2	BLTZAL	BGEZAL						
3	3								
4	4								
		COPz							
		0	1	2	3	4	5	6	7
25..24	23..21								
0	0	MF		CF		MT		CT	
1	1	BC	†	†	†	†	†	†	†
2	2								
3	3								
		Co-Processor Specific Operations							
		0	1	2	3	4	5	6	7
20..19	18..16								
0	0	BCzF	BCzT						
1	1								
2	2								
3	3								
		CPO							
		0	1	2	3	4	5	6	7
4..3	2..0								
0	0		TLBR	TLBWI				TLBWR	
1	1	TLBP							
2	2	RFE							
3	3								

Table 2.17 Opcode Encoding



Introduction

The R36100 achieves its high standard of performance by combining a fast, efficient execution engine (that of the R3000A) with high-memory bandwidth, supplied from its large internal instruction and data caches. These caches insure that the majority of processor execution occurs at the rate of one instruction per clock cycle, and serve to decouple the high-speed execution engine from slower, external memory resources.

Portions of this chapter review the fundamentals of general cache operation, and may be skipped by readers already familiar with these concepts. This chapter also discusses the particular organization of the on-chip caches of the R36100. However, as these caches are managed by the R36100 itself, the system designer does not typically need to be explicitly aware of this structure.

Fundamentals of Cache Operation

High-performance microprocessor-based systems frequently borrow from computer architecture principles long used in mini-computers and mainframes. These principles include instruction execution pipelining (discussed in Chapter 2) and instruction and data caching.

A cache is a high-speed memory store which contains the instructions and data most likely to be needed by the processor. That is, rather than implement the entire memory system with zero wait-state memory devices, a small zero wait-state memory is implemented. This memory, called a cache, contains the instructions/data most likely to be referenced by the processor. If indeed the processor issues a reference to an item contained in the cache, then a zero wait-state access is made; if the reference is not contained in the cache, then the longer latency associated with the true processor memory is incurred. The processor will achieve its maximum performance as long as its references "hit" (are resident) in the cache.

Caches rely on the principles of locality of software. These principles state that when a data/instruction element is used by a processor, it and its close neighbors are likely to be used again soon. The cache is then constructed to keep a copy of instructions and data referenced by the processor, so that subsequent references occur with zero wait-states.

Since the cache is typically many orders of magnitude smaller than main memory or the virtual address space, each cache element must contain both the data (or instruction) required by the processor, as well as information which can be used to determine whether a cache "hit" occurs. This information, called the cache "TAG", is typically some or all of the address in main memory of the data item contained in that cache element as well as a "Valid" flag for that cache element. Thus, when the processor issues an address for a reference, the cache controller compares the TAG with the processor address to determine whether a hit occurs.

To minimize cost while maintaining high-performance, the R36100 integrates a reasonable amount of cache internal to the chip, eliminating the cost and complexity of external caches.

R36100 Cache Organization

There are a number of algorithms possible for managing a processor cache. This section describes the cache organization and operation of the R36100.

Basic Cache Operation

When the processor makes a reference, its 32-bit internal physical address bus contains the address it desires. The processor address bus is split into two parts; the low-order address bits specify a location in the cache to access, and the remaining high-order address bits contain the value expected from the cache TAG. Thus, both the instruction/data element and the cache TAG are fetched simultaneously from the cache memory. If the value read from the TAG memories is the same as the high-order address bits, a cache hit occurs and the processor is allowed to operate on the instruction/data element retrieved. Otherwise, a cache miss is processed. This operation is illustrated in Figure 3.1.

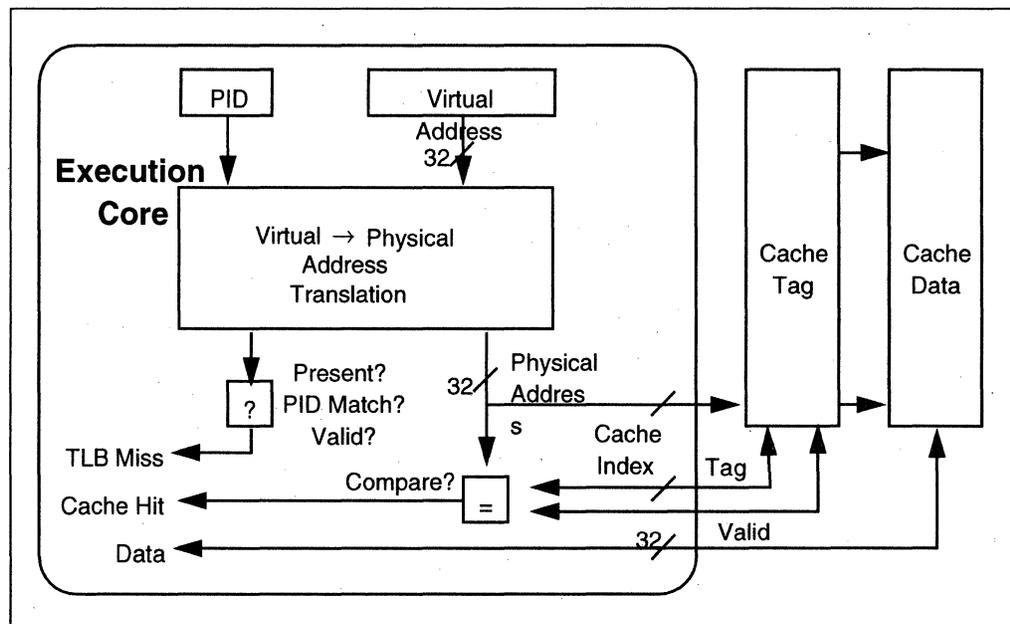


Figure 3.1 Cache Line Selection

To maximize performance, the R36100 implements a Harvard Architecture caching strategy. That is, there are two separate caches: one contains instructions (operations), and the other contains data (operands). By separating the caches, higher overall bandwidth to the execution core is achieved, and thus higher performance is realized.

Memory Address to Cache Location Mapping

The R36100's caches are direct-mapped. That is, each main memory address can be mapped to (contained in) only one particular cache location. This is different from set-associative mappings, where each main memory location has multiple candidate cache locations for address mapping.

This organization, coupled with the relatively large cache sizes resident on the R36100, achieve extremely high hit rates while maximizing speed and minimizing complexity and power consumption.

Cache Addressing

The address presented to the cache and cache controller is that of the physical (main) memory element to be accessed. That is, the virtual address to physical address translation is performed by the memory management unit prior to the processor issuing its reference address.

Some microprocessors utilize virtual indexing and virtual tagging in the cache, where the processor virtual address is used to specify the cache element to be retrieved. This type of cache structure complicates software and slows embedded applications:

- When the processor performs a context switch, a virtually tagged cache must be flushed. This is because two different tasks can use the same virtual address but mean totally different physical addresses. This cache flushing for a large cache dramatically slows context switch performance.
- Software must be aware of and specifically manage against “alias” problems. An alias occurs when two different virtual addresses correspond to the same physical address. If that occurs in a virtually indexed cache, then the same data element may be present in two different cache locations. If one virtual address is used to change the value of that memory location, and a different address used to read it later, then the second reference will not get the current value of that data item.

By providing for the virtual-to-physical address translation in the processor pipeline, physical cache addressing is used with no inherent performance penalty.

To support cache locking, the R36100 allows the kernel software to select certain high-order physical address bits to replace normal high-order cache index lines. This separates the cache into two portions: a lower portion, which services physical addresses below the high-order address; and a higher portion, which services physical addresses above the high-order address. Even when this mode is enabled, the R36100 implements direct-mapped, physically indexed, physically tagged caches.

Write Policy

The R36100 utilizes a write-through cache. That is, whenever the processor performs a write operation to memory, then both the cache (data and TAG fields) and main memory are written. If the reference is uncacheable, then only main memory is written.

To minimize the delays associated with updating main memory, the R36100 contains a 4 element write buffer. The write buffer captures the target address and data value in a single processor clock cycle, and subsequently performs the main memory write at its own, slower rate. The write buffer can FIFO up to 4 pending writes, as described in a later chapter.

Partial Word Writes

In the case of partial word writes (store operations of less than 32-bits), the R36100 operates by performing a read-modify-write sequence in the cache: the store target address is used to perform a cache fetch; if the cache “hits”, then the partial word data is merged with the cache and the cache is updated. If the cache read results in a hit, the memory interface will see the full word write, rather than the partial word. This allows the designer to observe the actual activity in the on-chip caches.

If the cache lookup of a partial word write “misses” in the cache, then only main memory is updated.

Instruction Cache Line Size

The "line size" of a cache refers to the number of cache elements mapped by a single TAG element. In the R36100, the instruction cache line size is 16 bytes, or four words.

This means that each cache line contains four adjacent words from main memory. In order to accommodate this, an instruction cache miss is processed by performing a quad word (block) read from the main memory, as discussed in a later chapter. This insures that a cache line contains four adjacent memory locations. Note that since the instruction cache is typically never written into directly by user software, the larger line size is permissible. If software does explicitly store into the instruction cache (perform store operations with the caches "swapped"), the programmer must insure that either the written lines are left invalidated, or that they contain four adjacent instructions.

Block refill uses the principle of locality of reference. Since instructions typically execute sequentially, there is a high probability that the instruction address immediately after the current instruction will be the next instruction. Block refill then brings into the cache those instructions immediately near the current instruction, resulting in a higher instruction cache hit rate.

Block refill also takes advantage of the difference between memory latency and memory bandwidth. Memory latency refers to the amount of time required to perform a processor request, while bandwidth refers to the rate at which subsequent transfers can occur. Factors that affect memory latency include address decoding, bus arbitration, and memory pre-charge requirements; factors which maximize bandwidth include the use of page mode or nibble mode accesses, memory interleaving, and burst memory devices.

The processing of a quad word read is discussed in a later chapter; however, it is worth noting that the R36100 can support either true "burst" accesses or can utilize a simpler, slower memory protocol for quad word reads. Also note that the variable bus sizing capability of the R36100 means that block reads can occur from 8- or 16-bit memory systems. This includes the case of instruction fetches; the bus interface unit will automatically translate the block read protocol into a larger number of sub-word reads, depending on the memory width programmed for the target memory location.

Finally, note that the R36100 performs "streaming" during instruction cache refill. That is, the processor will simultaneously refill the instruction cache and execute the incoming instructions. Streaming contributes an average of 5% of performance.

Data Cache Line Size

The data cache line size is different from that of the instruction cache, based on differences in their use. The data cache is organized as a line size of one word (four bytes).

This is optimal for the write policy of the data cache: since an individual cache word may be written by a software store instruction, the cache controller cannot guarantee that four adjacent words in the cache are from adjacent memory locations. Thus each word is individually tagged. The partial word writes (less than 4 bytes) are handled as a read-modify-write sequence, as described above.

Although the data cache line size is one word, the system may elect to perform data cache updates using quad word reads (block refill). The performance of the data cache update options can be measured in an actual system, by turning on the two different options under software control.

Summary

The on-chip caches of the R36100 family can be thought of as constructed from discrete devices around the R3000A. Figure 3.2 shows the block diagram of the cache interface for the R36100.

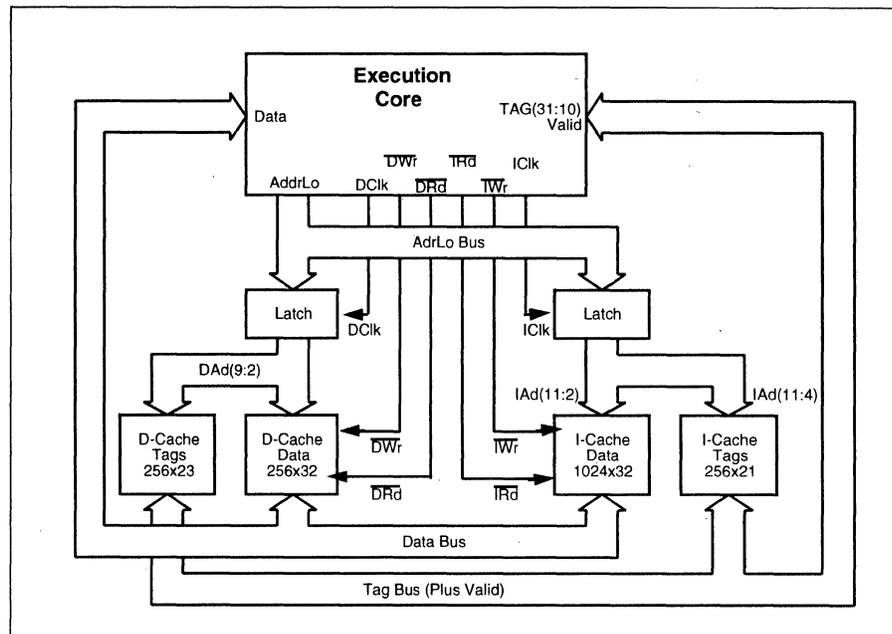


Figure 3.2 R36100 Execution Core and Cache Interface

Cache Operation

The operation of the on-chip caches is very straightforward, and is automatically handled by the processor.

Basic Cache Fetch Operation

As with the R3000A/R3500, the R36100 can access both the instruction and data caches in a single clock cycle, resulting in high bandwidth to the execution core. It does this by time multiplexing the cycle in the cache interface:

- During the first phase, a data cache address is presented, and a previous instruction cache read is completed.
- During the second phase, the data cache is read into the processor (or written by the processor). Also, the instruction cache is addressed with the next desired instruction.
- During the first phase of the next cycle, the instruction fetch begun in the previous phase is completed and a new data transaction is initiated.

This operation is illustrated in Figure 3.3 on page 6. As long as the processor hits in the cache, and no internal stall conditions are encountered, it will continue to execute *run* cycles. A run cycle is defined to be a clock cycle in which forward progress in the processor pipeline occurs. Note that data in the cache is organized into 32-bit words, regardless of the width associated with main-memory from which the datum was taken. Thus, cache hits can retrieve a full 32-bits in a single cycle, minimizing the performance impact of the narrower memory system.

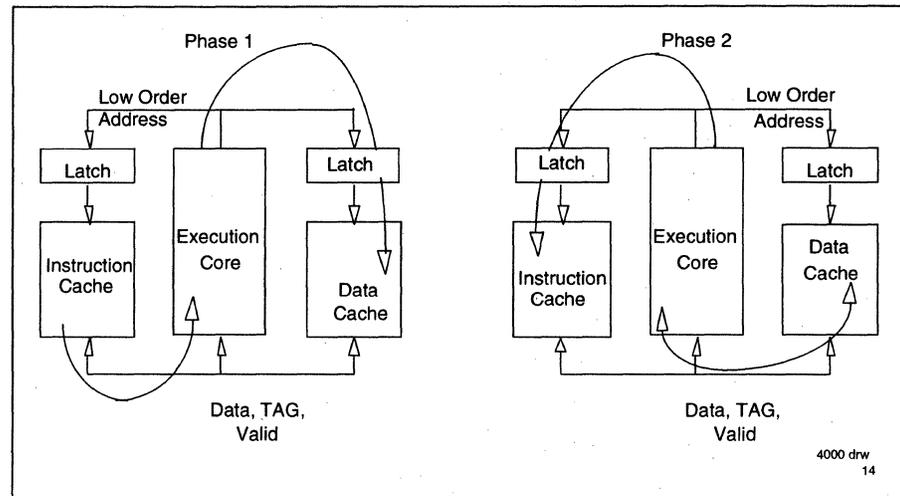


Figure 3.3 Phased Access of Instruction and Data Caches

Cache Miss Processing

In the case of a cache miss (due to either a failed tag comparison or because the processor issued an uncacheable reference), the main memory interface (discussed in a later chapter) is invoked. If, during a given clock cycle, both the instruction and data cache miss, the data reference will be resolved before the instruction cache miss is processed.

While the processor is waiting for a cache miss to be processed, it will enter *stall* cycles until the bus interface unit indicates that it has obtained the necessary data.

When the bus interface unit returns the data from main memory, it is simultaneously brought to the execution unit and written into the on-chip caches. This is performed in a processor *fixup* cycle.

During a fixup cycle, the processor re-issues the cache access that failed; this occurs by having the processor re-address the instruction and data caches, so that the data may be written into the caches and brought into the execution core. If the cache miss was due to an uncacheable reference, the write is not performed, although a fixup cycle does occur to allow the data to be brought into the execution core.

Instruction Streaming

A special feature of the R36100 is utilized when performing block reads for instruction cache misses. This process is called *instruction streaming*. Instruction streaming is simultaneous instruction execution and cache refill.

As the block is brought in, the processor refills the instruction cache. Execution of the instructions within the block begins when the instruction corresponding to the cache miss is returned by the bus interface unit to the execution core. Execution continues until the end of the block is reached (in which case normal execution is resumed), or until some event forces the processor core to discontinue execution of that stream. These events include:

- Taken branches
- Data cache miss
- Internal stalls (TLB miss, multiply/divide interlock)
- Exceptions

When one of these events occur, the processor re-enters simple cache refill until the rest of the block has been written into the cache, to insure that one TAG describes all four adjacent words.

Cacheable References

Chapter 4 explains how the processor determines whether a particular reference (either instruction or data) is to a memory location that may reside in the cache. The fundamental mechanism is that certain virtual addresses are considered to be “cacheable”. If the processor attempts to make a reference to a cacheable address, then it will employ its cache management protocol through that reference. Otherwise, the cache will be bypassed, and the execution engine core will directly communicate with the bus interface unit to process the reference.

Whether a given reference should be cacheable or not depends on the application and on the target of the reference. Generally, I/O devices should be referenced as uncacheable data; for example, if software was polling a status register, and that register was cached, then it would never see the device update the status (note that most compiler suites support the “volatile” data type to insure that the I/O device status register data in this case never gets allocated into an internal register). In the R36100, the cacheability of the on-chip registers in the I/O and peripheral devices is automatically selected to be “non-cacheable”.

There may be other instances where the uncacheable attribute is appropriate. For example, software which directly manipulates or flushes the caches can not be cached; similarly, boot software can not rely on the state of the caches, and thus must operate uncached at least until the caches are initialized.

Software Directed Cache Operations

In order to support certain system requirements, the R36100 provides mechanisms for software to explicitly manipulate the caches. These mechanisms support diagnostics, cache and memory sizing, and cache flushing. In general, these mechanisms are enabled/disabled through the use of the Status Register in CPO.

The primary mechanisms for supporting these operations are cache swapping and cache isolation. Cache swapping forces the processor to use the data cache as an instruction cache, and vice versa. It is useful for allowing the processor to issue store instructions which cause the instruction cache to be written. Cache isolation causes the current data cache to be “isolated” from main memory; store operations do not cause main memory to be written, and all load operations “hit” in the data cache.

These mechanisms are enabled through the use of the “IsC” (Isolate Cache) and “SwC” (Swap Cache) bits of the status register, which resides in the on-chip System Control Co-Processor (CPO). *The 5 instructions which immediately precede and succeed these operations must not be cacheable, so that the actual swapping/isolation of the cache does not disrupt operation.*

Cache Sizing

It is possible for software to determine the amount of cache resident on any given R3xxx-based chip (note that the R3041, R3051, R3052, and R3071/R3081 each feature differing amounts of cache on chip). Having software determine the size of the cache at boot time, rather than building static values into the software, allows for maximum flexibility in using various members of the R3xxx family, including future devices.

Cache sizing in an R36100 is performed much like traditional memory sizing algorithms, but with the cache isolated. This avoids side-effects in memory from the sizing algorithm, and allows the software to use the “Cache Miss” bit of the status register in the sizing algorithm.

To determine the size of the instruction cache, software should:

1. Swap Caches (not needed for D-Cache sizing)
2. Isolate Caches
3. Write a value at location 8000_0000
4. Write a value at location 8000_0200 (8000_0000 + 512B)

Read location 8000_0000.

Examine the CM (Cache_Miss) bit of the status register; if it indicates a cache miss, then the cache is 512B; otherwise, the cache is 1kB or larger.

5. Write a value at location 8000_0400 (8000_0000 + 1kB)

Read location 8000_0000.

Examine the CM (Cache_Miss) bit of the status register; if it indicates a cache miss, then the cache is 1KB; otherwise, the cache is 2KB or larger.

6. etc...

Of course a more generalized algorithm could be developed to determine the cache size; this may be desirable for compatibility with discrete R3000A/R3500 systems or other R3051 family members. However, any algorithm will probably include the Swap and Isolate of the Instruction Cache, and the use of the Cache Miss bit. Sizing the data cache is done with a similar algorithm, although the caches need not be swapped, and smaller cache sizes need to be considered.

Note that this software should operate as uncached. Once this algorithm is done, software should return the caches to their normal state by performing either a complete cache flush or an invalidate of those cache lines modified by the sizing algorithm.

Cache Flushing

Cache flushing refers to the act of invalidating (indicating a line does not have valid contents) lines within either the instruction or data caches. Flushing must be performed before the caches are first used as real caches, and might also be performed during main memory page swapping or at certain context switches (note that the R3051 family implements physically addressed caches, so that cache flushing at context switch time is not generally required).

The basic concept behind cache flushing is to have the "Valid" bit of each cache line set to indicate invalid. This is done in the R36100 by having the cache isolated, and then writing a partial word quantity into the current data cache. Under these conditions, the CPU will negate the "Valid" bit of the target cache line.

Again, this software should operate as uncached. To flush the data cache:

1. Isolate Caches
2. Perform a byte write every 4 bytes, starting at location 0, until 256 such writes have been performed (128 in the R3041, more for other R3xxx family members).
3. Return the data cache to its normal state by clearing the IsC bit.

To flush the instruction cache:

1. Swap Caches
2. Isolate Caches
3. Perform a byte write every 16 bytes (based on the instruction cache line size of 16 bytes). This should be done until each line (256 lines in the R36100, more or less for other R3xxx devices) have been invalidated. Note that treating the R36100 as if it had larger on-chip caches, and flushing/invalidating more than 256 lines is acceptable though less efficient.
4. Return the caches to their normal state (unswapped and not isolated).

To minimize the execution time of the cache flush, this software should probably use an “unrolled” loop. That is, rather than have one iteration of the loop invalidate only one cache line, each iteration should invalidate multiple lines. This spreads the overhead of the loop flow control over more cache line invalidates, thus reducing execution time.

Also, of course it is preferable to use the cache sizing algorithm described earlier to determine the number of lines to be flushed.

Forcing Data into the Caches

Using these basic tools, it is possible to have software directly place values into the caches. When combined with appropriate software techniques, this could be used to “lock” values into the on-chip caches, by insuring that software does not issue other cacheable address references which may displace these locked values.

In order to force values into a cache, the cache should be Isolated. If software is trying to write instructions into the instruction cache, then the caches should also be swapped.

When forcing values into the instruction cache, software must take care with regards to the line size of the instruction cache. Specifically, a single TAG and Valid field describe four words in the instruction cache; software must then insure that any instruction cache line tagged as Valid actually contains valid data from all four words of the block.

Cache-Locking Operation

The R36100 implements the ability to segregate the caches into 2 or 4 portions, or to allow it to operate as a normal single contiguous entity. Either or both the instruction and data cache can be run in any split or non-split mode independently.

As an example, splitting the cache into halves or quarters allows interrupt service routines and data to be locked into part of the cache, while the remainder of the cache is used for the user program and data.

If run in the normal mode (as a single contiguous entity), the cache index (used internally to address the Cache Data and Tag RAMs) is derived solely from the low-order physical address bits. For example, the cache index for the data cache is PhysAddr(9:2); and for the instruction cache, the cache index is PhysAddr(11:2).

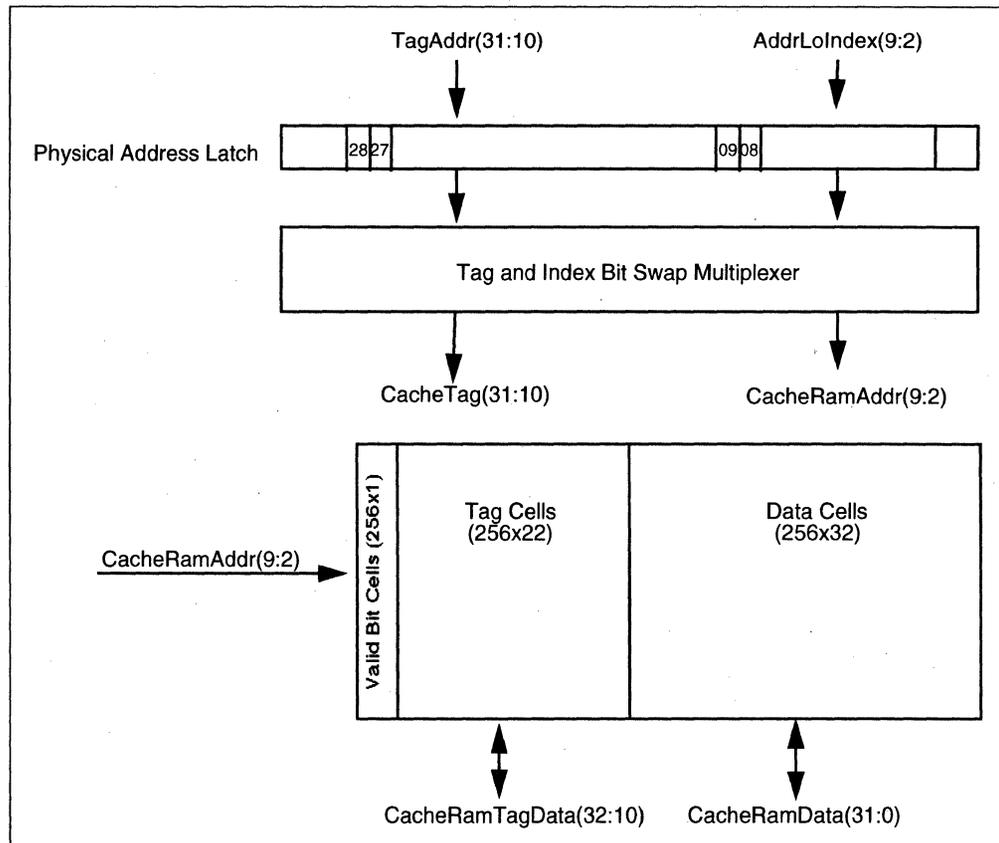


Figure 3.4 R36100 Instruction Cache Index Address Path

In the normal mode case, a reference with the same low-order PhysAddr bits but different high-order PhysAddr tags will cause the current cache contents to be replaced. For example, location 0x0000_1008 will be entered into the line at cache index 0x0000; if that line previously was cached with main memory location 0x0000_0008, it would be replaced with new data and tag. Any address which is modulo 4kB (for instance 0x1004_0008) could cause replacement of that cache line.

On the other hand, in the split modes, the system software can instruct the cache controller to use either or both of PhysAddr(28:27) as the uppermost two index bits (2 or 4 portions). In this case, the cache simultaneously direct-maps multiple distinctly different memory spaces. The 10 bits for the instruction cache index can be constructed as PhysAddr(28, 10:2), for example.

When the caches are operated in split mode, typically the MSBs going to main memory (bits 31-29) are masked out from the physical address decode. On the R36100, the physical address decode is part of either the DRAM Controller's or the Memory/IO Controller's Page Register and Page Mask Register (whichever one is used to control main memory). Masking out the MSBs allows physical RAM space to be contiguous (for instance contained within a 1MB block), while the virtual program space can vary the MSBs. For instance the virtual program space can consist of a 512KB block beginning at 0x0000_0000 and a second virtual program space with a 512KB block beginning at 0x1004_0000. The two virtual addresses will translate (see the next chapter for more details) to physical addresses 0x4000_0000 and 0x5004_0000, respectively, as far as cache memory is concerned. Since bits 31-29 are ignored by the main memory controller, the two physical addresses are effectively 0x0000_0000 and 0x0004_0000 as far as main memory RAM is concerned. Thus by using the Page Mask Register, the caches can see 2 or 4 blocks of address spaces, while main memory sees a single large block of address space.

To continue the instruction cache example, the upper 2kB portion of the I-cache services physical addresses in the range of 0x1000_0000 and above; physical addresses in the range 0x0fff_ffff and below are serviced by the lower 2kB I-cache portion.

In this example, the instruction at physical location 0x1004_0008 will not replace the contents of the line which holds memory location 0x0000_0008. These two portions of software will not interfere with each other in the caches. The software developer typically specifies the address region for code in either the kernel, or with the linker. This mechanism allows the programmer to separate code into portions and independently lock them, without requiring page management software or complex operating system software.

Physical address 0x0000_0008 is accessed via kuseg (explained in the next chapter), and is typically in the area of the exception vector. Physical address 0x1000_0008 is spaced 256MB higher in memory. In this system, system tasks operating higher in kuseg or in kseg0 do not "knock out" the exception service code, effectively locking this time critical code into one half of the on-chip cache.

Table 3.1 on page 13 shows the correlation between physical address lines, cache index lines, and cache sub-segments supported by the R36100. Figure 3.5, Figure 3.6, and Figure 3.7 on page 12 shows the mapping of physical addresses to cache when the cache is 1, 2, or 4 portions.

Note that these tables and drawings assume that the code operates out of kuseg (explained in the next chapter). Since the R36100 implements 32-bit virtual and physical addressing, the patterns shown repeat every time a very high-order (PhysAddr(29) and above) is changed; thus, there are 8 such copies of each cache region, separated by 512MB each. The tables and example assume that PhysAddr(31:29) are all '0' throughout system software. However, memory spaces larger than 512MB are rarely used with embedded systems, the example in these tables will suffice for almost all systems.

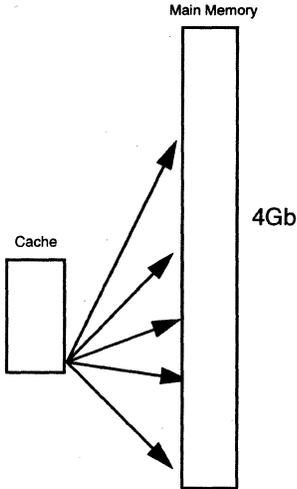


Figure 3.5 R36100 Cache in One Portion

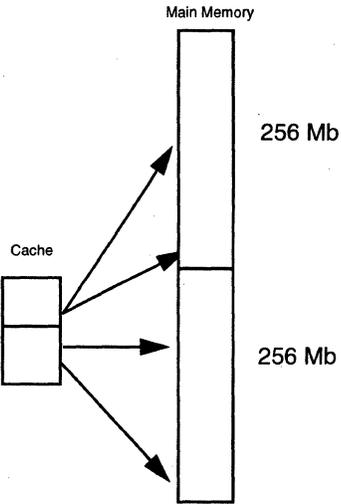


Figure 3.6 R36100 Cache in Two Portions

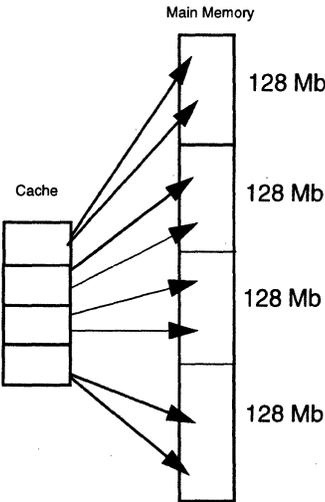


Figure 3.7 R36100 Cache in Four Portions

Cache IndexAddr(11)	Cache IndexAddr(10)	Physical Address Range	Cache Size
PhysAddr(11)	PhysAddr(10)	0x0000_0000 0x1FFF_FFFF	4kB
PhysAddr(28)	PhysAddr(10)	0x0000_0000 - 0x0FFF_FFFF	2kB
		0x1000_0000 - 0x1FFF_FFFF	2kB
PhysAddr(28)	PhysAddr(27)	0x0000_0000 - 0x07FF_FFFF	1kB
		0x0800_0000 - 0x0AFF_FFFF	1kB
		0x1000_0000 - 0x17FF_FFFF	1kB
		0x1800_0000 - 0x1FFF_FFFF	1kB
Note: This table describes byte-addressable caches, with the lsb of the cache index == 2.			

Table 3.1 Instruction Cache to Address Mapping under Various Cache Locking Conditions

Cache Index Addr(9)	Cache Index Addr(8)	Physical Address Range	Cache Size
PhysAddr(9)	PhysAddr(8)	0x0000_0000 - 0x1FFF_FFFF	1kB
PhysAddr(28)	PhysAddr(8)	0x0000_0000 - 0x0FFF_FFFF	512B
		0x1000_0000 - 0x1FFF_FFFF	512B
PhysAddr(28)	PhysAddr(27)	0x0000_0000 - 0x07FF_FFFF	256B
		0x0800_0000 - 0x0AFF_FFFF	256B
		0x1000_0000 - 0x17FF_FFFF	256B
		0x1800_0000 - 0x1FFF_FFFF	256B
Note: This table describes byte-addressable caches, with the lsb of the cache index == 2.			

Table 3.2 Data Cache to Address Mapping under Various Cache Locking Conditions

Summary

The on-chip caches of the R36100 are key to the inherent performance of the processor. The R36100 design, however, does not require the system designer (either software or hardware) to explicitly manage this important resource, other than to correctly choose virtual addresses which may or may not be cached, and to flush the caches at system boot. This contributes to both the simplicity and performance of an R36100 system.



Integrated Device Technology, Inc.

Virtual-to-Physical Address Translation and Address Map

Chapter 4

The R36100 provides the same basic virtual-to-physical address translation as the rest of the R30xx family base versions (the R3041, R3051, R3052, R3071, and R3081). These devices provide segment-based virtual-to-physical address translation, and support the segregation of kernel and user tasks without requiring extensive virtual page management.

The extended versions of the R30xx family (the R3051E, R3052E, R3071E, and R3081E) provide a full featured memory management unit (MMU) identical to the MMU structure of the R3000A and R3500. The extended MMU uses an on-chip translation lookaside buffer (TLB) and dedicated registers in CP0 to provide for software management of page tables. There is no Extended Architecture version of the R36100.

This chapter describes the operating states of the processor (kernel and user), and describes the virtual-to-physical address translation mechanisms provided in the R36100.

Virtual Memory in the R3000A Architecture

There are two primary purposes of the memory management capabilities of the R3000A Architecture:

- Various areas of main memory can have individual sets of attributes associated with them. For example, some segments may be indicated as requiring kernel status to be accessed; others may have cacheable or uncacheable attributes. The virtual-to-physical address translation establishes the rules appropriate for a given virtual address. The R36100 memory manager provides for these mechanisms, without requiring the use of a TLB.
- The virtual memory system can be used to logically expand the physical memory space of the processor, by translating addresses composed in a large virtual address space into the physical address space of the system. This is particularly important in applications where software may not be explicitly aware of the hardware resources of the processor system, and includes applications such as X-Window display systems. These types of applications may be better served by the "E" (extended architecture) versions of the R30xx family. On the other hand, certain real-time operating systems offer similar functionality without requiring an MMU; for example, the IDT/c tool chain supports position-independent code without requiring a page fault manager in the operating system.

Figure 4.1 shows the format of an R3000A architecture virtual address. The most significant 20 bits of the 32-bit virtual address are called the virtual page number, or VPN. In the extended architecture versions, the VPN allows mapping of virtual addresses based on 4kB pages; in the base versions (and thus in the R36100), only the three highest bits (segment number) are involved in the virtual-to-physical address translation.

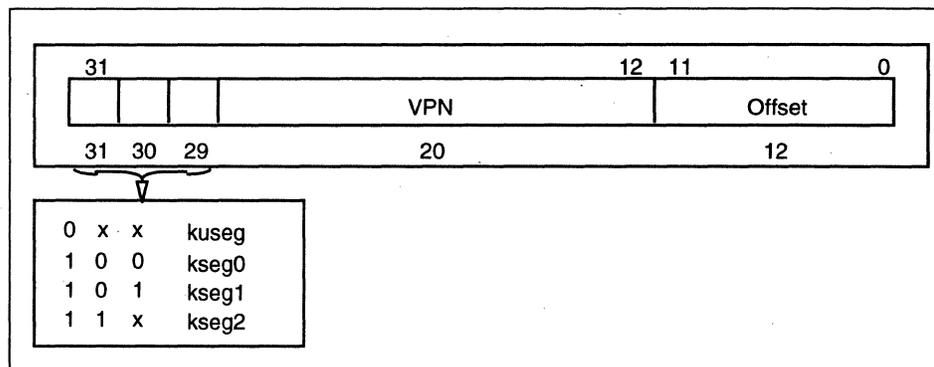


Figure 4.1 Virtual Address Format

The three most significant bits of the virtual address identify which virtual address segment the processor is currently referencing; these segments have associated with them the mapping algorithm to be employed, and whether virtual addresses in that segment may reside in the cache. The translation of the virtual address to an equivalent privilege level/segment is the same for the base and extended versions of the architecture.

Privilege States

The R36100 provides for two unique privilege states: the “Kernel” mode, which is analogous to the “supervisory” mode provided in many systems, and the “User” mode, where non-supervisory programs are executed. Kernel mode is entered whenever the processor detects an exception; when a Restore From Exception (RFE) instruction is executed, the processor will return either to its previous privilege mode or to User mode, depending on the state of the machine and when the exception was detected.

User Mode Virtual Addressing

While the processor is operating in User mode, a single, uniform virtual address space (**kuseg**) of 2GB is available for Users. All valid user-mode virtual addresses have the most significant bit of the virtual address cleared to 0. An attempt to reference a Kernel address (most significant bit of the virtual address set to 1) while in User mode will cause an Address Error Exception. Kuseg begins at virtual address 0 and extends linearly for 2GB. This segment is typically used to hold user code and data, and the current user processes.

Also note that the physical address space corresponding to kuseg is independent of the physical address spaces of the various kernel only segments. Thus, systems can be constructed which preclude user tasks from affecting kernel memory. On the other hand, simple systems can, by virtue of the address decode, compress the mapping into a single address region.

Kernel Mode Virtual Addressing

When the processor is operating in Kernel mode, four distinct virtual address segments are simultaneously available. The segments are:

- **kuseg.** The kernel may assert the same virtual address as a user process, and have the same virtual-to-physical address translation performed for it as the translation for the user task. This facilitates the kernel having direct access to user memory regions. The virtual-to-physical address translation, including the Port Size attributes, is identical with User mode addressing to this segment.

- **kseg0.** Kseg0 is a 512MB segment, beginning at virtual address 0x8000_0000. This segment is always translated to a linear 512MB region of the physical address space starting at physical address 0. All references through this segment are cacheable.

When the most significant three bits of the virtual address are “100”, the virtual address resides in kseg0. The physical address is constructed by replacing these three bits of the virtual address with the value “000”. As these references are cacheable, kseg0 is typically used for kernel executable code and some kernel data.

- **kseg1.** Kseg1 is also a 512MB segment, beginning at virtual address 0xa000_0000. This segment is also translated directly to the 512MB physical address space starting at address 0. All references through this segment are uncacheable.

When the most significant three bits of the virtual address are “101”, the virtual address resides in kseg1. The physical address is constructed by replacing these three bits of the virtual address with the value “000”. Unlike kseg0, references through kseg1 are not cacheable. This segment is typically used for I/O registers, boot ROM code, and operating system data areas such as disk buffers.

- **kseg2.** This segment is analogous to kuseg, but is accessible only from kernel mode. This segment contains 1GB of linear addresses, beginning at virtual address 0xc000_0000. As with kuseg, the virtual-to-physical address translation depends on whether the processor is a base or extended architecture version.

When the two most significant bits of the virtual address are “11,” the virtual address resides in the 1024MB segment kseg2. The virtual-to-physical translation is done either through the TLB (extended versions of the processor) or through a direct segment mapping (base versions). An operating system would typically use this segment for stacks, per-process data that must be re-mapped at context switch, user page tables, and for some dynamically allocated data areas.

Base versions of the R30xx family (including the R36100) are distinguishable from extended versions in software by examining the TS (TLB Shutdown) bit of the Status Register after reset, before the TLB is used. If the TS bit is set (1) immediately after reset, indicating that the TLB is non-functional, then the current processor is a base version of the architecture. If the TS bit is cleared after reset, then the software is executing on an extended architecture version of the processor.

The PRId register—described in a later chapter—can be used to distinguish the R36100 from other members of the R30xx family.

R36100 address translation

Processors which only implement the base versions of memory management perform direct segment mapping of virtual-to-physical addresses, as illustrated in Figure 4.2. Thus, the mapping of kuseg and kseg2 is performed as follows:

- Kuseg is always translated to a contiguous 2GB region of the physical address space, beginning at location 0x4000_0000. That is, the value "00" in the two highest order bits of the virtual address space are translated to the value "01", and "01" is translated to "10", with the remaining 30 bits of the virtual address unchanged.
- Virtual addresses in kseg2 are directly output as physical addresses; that is, references to kseg2 occur with the physical address unchanged from the virtual address.
- Virtual addresses in kseg0 and kseg1 are both translated identically to the same physical address region.

The base versions of the architecture allow kernel software to be protected from user mode accesses, without requiring virtual page management software. User references to kernel virtual address will result in an address error exception.

Note that the special areas of the virtual address space shown in Figure 4.2 are translated to physical addresses identically with the remainder of their virtual address segment. In the R30xx family, these address areas were indicated as "reserved" for compatibility with future devices.

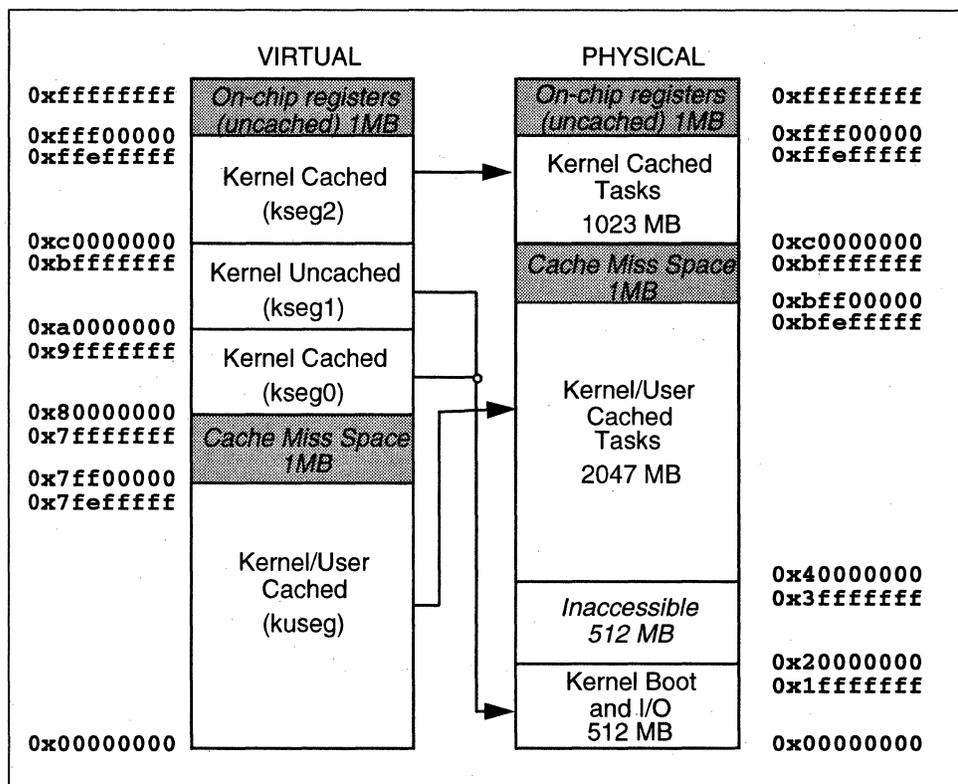


Figure 4.2 virtual-to-physical Address Translation in R36100

Some systems may elect to protect external physical memory as well. That is, the system may include distinct memory devices which can only be accessed from kernel mode. The physical address output determines whether the reference occurred from kernel or user mode, according to Table 4.1. Some systems may wish to limit accesses to some memory or I/O devices to those physical address bits which correspond to kernel mode virtual addresses.

Alternately, some systems may wish to have the kernel and user tasks share common areas of memory. Those systems could choose to have their address decoder ignore the high-order physical address bits, and compress all of memory into the lower region of physical memory. The high-order physical address bits may be useful as privilege mode status outputs in these systems.

Physical Address (31:29)	Virtual Address Segment
'000'	Kseg0 or Kseg1
'001'	Inaccessible
'01x'	Kuseg
'10x'	Kuseg
'11x'	Kseg2

Table 4.1 Virtual and Physical Address Relationships in Base Versions

On-Chip Registers

The top 1MB of virtual memory—which resides in the protected kernel space, `kseg2`—is treated as “non-cacheable” by the cache controller. The rest of `kseg2` is treated as cacheable. The on-chip memory controllers and peripherals have their register sets mapped into this address space; these registers need to be uncached to insure proper operation. Table 4.2 shows the address map for the on-chip resources.

Note that writes to addresses above `0xFFFF_E000` are propagated out to the external bus. However, none of the memory controllers are activated. This feature is provided to facilitate debug and in-circuit emulation equipment. Reads in this address range are propagated to the external bus.

Base Virtual Address	On-chip Resource
0xFFFF 8000	External Debug/Emulator Controller
0xFFFF 9000	Reserved
0xFFFF A000	
0xFFFF B000	
0xFFFF C000	
0xFFFF D000	
0xFFFF E000	
0xFFFF E100	DRAM Controller
0xFFFF E200	Memory and IO Controller
0xFFFF E300	Internal DMA Controller
0xFFFF E400	External DMA Controller
0xFFFF E500	Internal Debug/Emulator Controller
0xFFFF E600	Reserved
0xFFFF E700	Reserved
0xFFFF E800	Serial Port Interface
0xFFFF E900	Timer Interface
0xFFFF EA00	PIO Interface
0xFFFF EB00	Interrupt Peripheral Interface
0xFFFF EC00	Centronics Interface (P1284 interface)
0xFFFF ED00	Laser Printer Engine Interface
0xFFFF EE00	
0xFFFF EF00	
0xFFFF F000	

Table 4.2 R36100 On-Chip Resources and Address Map

As a general rule, the registers residing above `0xFFFF_E000` are 16-bits and in some cases 8-bits wide. Thus these registers require either half-word or byte load and store instructions for proper access. Because of the less-than-a-word access, if the system is big endian, the registers will either need a halfword offset of `0x2` or a byte offset of `0x3`. Little endian systems do not need an offset.

Cache Miss Area

The top 1MB of kuseg is also special. In the R36100, this area is the “Cache Miss” area.

If software attempts to “load” data with a modulo 16 address (lowest 4 address bits == 0), the cache controller will consider the access to have “missed” in the cache, regardless of the current tag contents.

This operation can speed certain types of data movement operations, especially when the contents of the corresponding main memory area may be updated externally to the processor. For example (See Table 4.3) if the main memory is a FIFO type memory, the code may perform a load to the FIFO address; the memory controller would burst four words into the cache (presuming a data block refill setting of four words) and load word “0” into the target register. The remaining words of the quad word read would be accessed from the cache. Once all four words are consumed, the code would issue another load with an offset of “0”, causing another cache miss process to the FIFO. Burst data movement is faster, since the software does not need to explicitly flush the cache line between bursts, nor does it need to use slower “uncached” single datum transfers.

```
#define FIFO_BASE 0x7FF00000      /* phys addr is 0xBFF00000 */

get_fifo:
li          t0, FIFO_BASE
lw          t1, 0x00(t0)
lw          t2, 0x04(t0)
lw          t3, 0x08(t0)
lw          t4, 0x0C(t0)          /* 13 cached clocks per 4 words */
```

Table 4.3 Example: FIFO load code using FCM memory space.

Summary

The R30xx family architecture provides two models of memory management: a very simple, segment based mapping, found in the base versions of the architecture, and a more sophisticated, TLB-based page mapping scheme, present in the extended versions of the architecture. Each scheme has advantages to different applications. The R36100 only implements the base version address translation in order to support low-cost systems.



Introduction

The MIPS architecture separates a processor into two (or three, in the case of a device with on-chip FPA) functional units: one is the general purpose CPU, which executes the actual code, and remains compatible across all of the devices; the other portion is the system control coprocessor (CPO) which manages the machine state, the virtual to physical address translation, and other device specific attributes.

This separation allows devices to be tailored to the needs of specific applications (by modifying CPO), yet retain software compatibility for the actual application itself (via the compatible CPU).

This chapter describes the implementation of CPO found on the R36100. In general, the exception handling methods of the R36100 are identical with those of the rest of the R30xx family; the memory management resources are identical with those of the base versions of the R30xx family; the only significant difference between this device and the R30xx family is in the implementation of the Cache Control register.

Coprocessor 0 Bus Interface Control

Figure 5.1 illustrates the coprocessor 0 registers found in the R36100. Note that the MIPS architecture allows the register set of CPO to vary by implementation; software can easily identify the R36100 (and its CPO registers) from other devices by reading the PRID from CPO.

The fields of these registers are described below. Table 5.1 lists the register numbers for the various R36100 CPO registers.

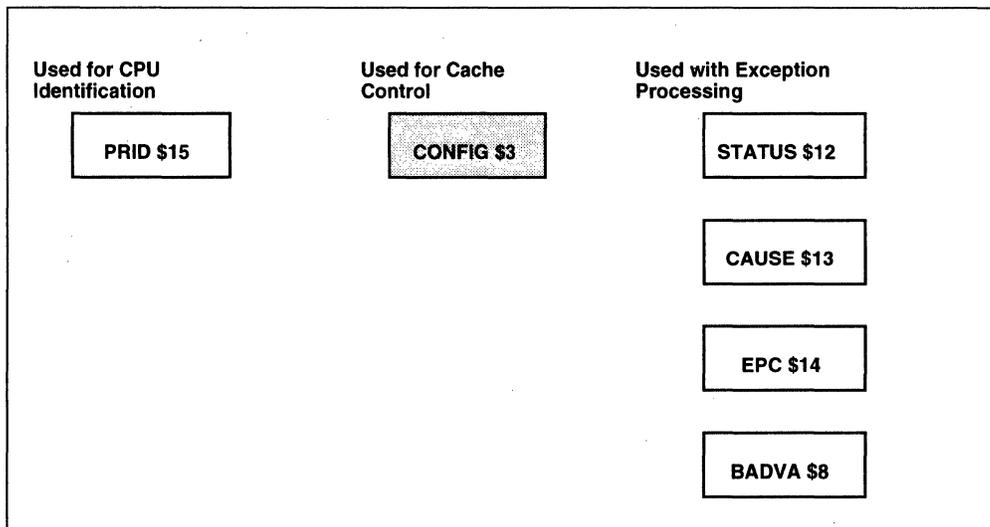


Figure 5.1 R36100 CPO Registers

Mnemonic	CPO Register #	Description
Config	\$3	Cache and CPU configuration control
BadVA	\$8	Bad Virtual Address for last Addressing Exception
Status	\$12	Processor status, control, and diagnostic information
Cause	\$13	Cause of current exception/exception state
EPC	\$14	Exception Program Counter; return address for exception handler
PrID	\$15	Identification information for current processor

Table 5.1 R36100 CPO Register Addresses

Cache Configuration Register

The cache configuration register allows the kernel to control various operational aspects of the on-chip caches of the R36100. These features can be used to improve performance and/or implement debug capability for the R36100. The Config register is both readable and writable.

Figure 5.2 illustrates the various fields of the cache configuration register. The reset defaults for this register insure R30xx compatible operation.

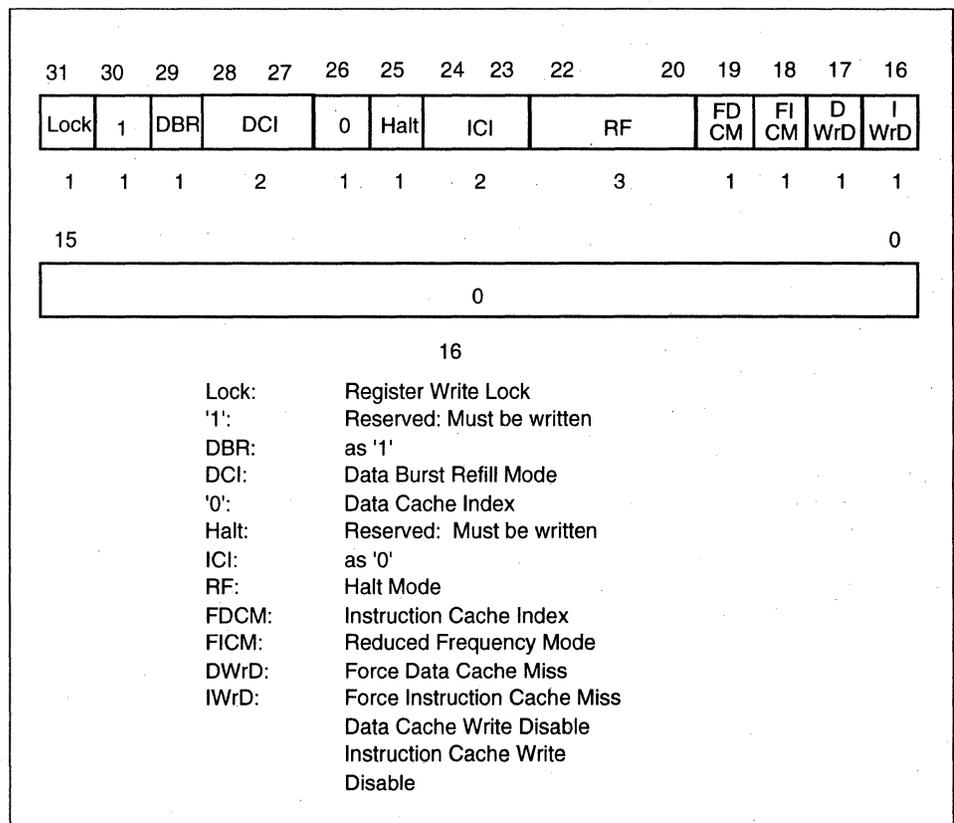


Figure 5.2 R36100 Cache Control Register

Lock ('Lock')

The lock bit can be used by the kernel to inhibit subsequent write operations to this register. It is useful in ensuring that operating systems written for other R30xx-based applications do not inadvertently change the fields of the Cache Configuration register.

Table 5.2 illustrates the Cache Configuration Register Lock Field. At reset, the register is unlocked (Lock bit is '0'). Thus, the Config register can be written and re-written as the operating system chooses. Once the Lock bit is written with a '1', subsequent writes to the Config register will be ignored.

Value	Action
'0'	Leave Unlocked (Default)
'1'	Lock register from future writes

Table 5.2 R36100 Cache Configuration Register Lock Field

Reserved-High ('1')

This bit is reserved for testing of the R36100. At reset, the bit will be set high ('1'). Writes to the Config register must maintain this bit as high ('1').

Reserved-Low ('0')

These fields are reserved for testing and for future R3xxx-based devices. At reset, these bit fields are reset ('0'). Writes to the Config register must maintain these bit fields as low ('0').

DBlockRefill ('DBR')

Table 5.3 indicates the value and action of the DBR. If this bit is set high ('1'), data cache misses will be processed as a quad (four-word) read. If this bit is reset low ('0'), data cache misses will be processed as a single word read. At reset, this bit is reset low ('0').

Value	Action
'0'	Data cache misses use single word refill (default).
'1'	Data cache misses use quad word refill.

Table 5.3 R36100 DBlockRefill Field

D-CacheIndexControl ('DCI')

This two bit field controls which bits of the physical address provide the high-order data cache index, as described in Chapter 3. Table 5.4 shows the actions of the various bit combinations. At reset, this field is cleared to '00', resulting in normal operation.

Value	DCache Index(9)	DCache Index(8)	Cache Portions
'00'	PhyAddr(9)	PhyAddr(8)	1 (default)
'01'	reserved	reserved	Not useful. Reserved
'10'	PhyAddr(28)	PhyAddr(8)	2
'11'	PhyAddr(28)	PhyAddr(27)	4

Table 5.4 R36100 D-Cache Index Control Field

Halt Mode ('Halt')

If this bit is set high ('1'), the CPU pipeline will be stalled until either an interrupt is asserted (regardless of current masking) or a reset exception is signalled. If this bit is set low ('0'), the pipeline will continue operation.

If the halt mode is exited, for example by an interrupt, the RF mode (described below) will also be exited. Table 5.5 shows the actions and values of the R36100 Halt Mode ('Halt').

Value	Action
'0'	Normal pipeline operation (default).
'1'	Halt until interrupt or Reset

Table 5.5 R36100 Halt Field

I-CacheIndexControl ('ICF')

This two bit field controls which bits of the physical address provide the high-order instruction cache index, as described in chapter 3. Table 5.6 shows the actions of the various bit combinations. At reset, this field is cleared to '00', resulting in normal operation.

Value	ICache Index(11)	ICache Index(10)	Cache Portions
'00'	PhyAddr(11)	PhyAddr(10)	1 (default)
'01'	reserved	reserved	Not useful. Reserved
'10'	PhyAddr(28)	PhyAddr(10)	2
'11'	PhyAddr(28)	PhyAddr(27)	4

Table 5.6 R36100 I-Cache Index Control Field

ReduceFrequency ('RF')

This 3 bit field can be used to divide the normal pipeline frequency down to a lower frequency, thus lowering device power consumption. Table 5.7 shows the actions of the various bit settings. At reset, this field is cleared to '000', resulting in normal operation. Similarly, whenever the halt mode is exited, this field will be cleared to '000'.

Value	Action
'000'	Normal Pipeline frequency (default)
'001'	Divide by 2
'010'	Divide by 4
'011'	Divide by 8
'100'	Divide by 16
'101'	Divide by 32
'110'	Divide by 64
'111'	Reserved

Table 5.7 R36100 Reduced Frequency Mode Field

When a reduced frequency mode is enabled, both the pipeline frequency and the system interface frequency will be reduced by the programmed amount. The minimum allowed frequency is a CPU pipeline frequency of 0.5MHz. To prevent internal synchronization problems, software should always switch from the Normal frequency to a particular divide by frequency or vice-versa. Thus if a switch between 64 and 32 is desired, first switch from 64 to Normal and then to 32.

Note that the "RF" mode also impacts the frequency of the bus interface, including the on-chip devices. System software may need to adjust timer values, baud rates, DRAM refresh, and other frequency sensitive system variables when entering and exiting "RF" mode.

ForceDCacheMiss ('FDCM')

Table 5.8 shows the values and actions for the R36100 ForceDCacheMiss field. If this bit is set high ('1'), all cacheable data load references will be forced to miss in the data cache. The data references will then be supplied using the Data Cache miss protocol (including DBlockRefill). Store operations will continue to update the cache, and the cache miss processing will update the cache. Thus, this bit provides a quick method of initializing the cache or reloading the cache from an external device.

At reset, this bit is reset low ('0'), allowing normal operation of the data cache. Note also that this bit is logically "OR'ed" with the emulator interface "FCM" pin.

Value	Action
'0'	Normal data cache operation (default).
'1'	Force data cache operations to miss.

Table 5.8 R36100 ForceDCacheMiss Field

ForceICacheMiss ('FICM')

Table 5.9 shows the values and actions for the ForceICacheMiss field. If this bit is set high ('1'), all cacheable instruction references will be forced to miss in the instruction cache. The instruction references will then be supplied using the Instruction Cache miss protocol (a quad word read). Cache miss processing will update the cache. Thus, this bit provides a quick method of initializing the cache or reloading the cache from an external device.

At reset, this bit is reset low ('0'), allowing normal operation of the instruction cache. Note also that this bit is logically "OR'ed" with the emulator interface "FCM" pin.

Value	Action
'0'	Normal instruction cache operation (default).
'1'	Force instruction cache operations to miss.

Table 5.9 R36100 ForceICacheMiss Field

DCacheWriteDisable('DWrD')

Table 5.10 shows the values and actions for the Data Cache Write Disable field. According to this table, when set high ('1'), this field causes data cache writes to be ignored. The data cache will thus contain the older value, regardless of the reason for the cache miss processing. Similarly, store instructions will not cause the D-cache to be updated. When cleared low ('0'), normal cache operation results.

Value	Action
'0'	Normal data cache operation (default).
'1'	Data cache writes inhibited.

Table 5.10 R36100 Data Cache Write Disable Field

I-CacheWriteDisable ('IWrd')

Table 5.11 shows the values and actions for the Instruction Cache Write Disable field. According to this table, when set high ('1'), this field causes instruction cache writes to be ignored. The instruction cache will thus contain the older value, regardless of the reason for the cache miss processing. When cleared low ('0'), normal cache operation results.

Value	Action
'0'	Normal data cache operation (default).
'1'	Instruction cache writes inhibited.

Table 5.11 R36100 Instruction Cache Write Disable Field

The Cause Register

The contents of the Cause register describe the last exception. A 5-bit exception code indicates the cause of the current exception; the remaining fields contain detailed information specific to certain exceptions.

All bits in this register, with the exception of the SW bits, are read-only. The SW bits can be written to set or reset software interrupts. Figure 5.3 illustrates the format of the Cause register. Table 5.12 details the meaning of the various exception codes.

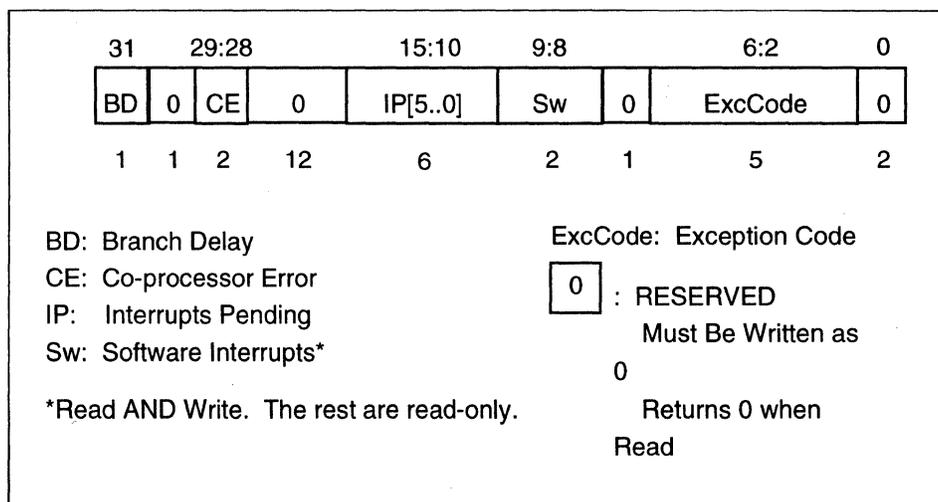


Figure 5.3 R36100 Cause Register

Number	Mnemonic	Description
0	Int	External Interrupt
1	MOD†	TLB Modification Exception
2	TLBL†	TLB miss Exception (Load or instruction fetch)
3	TLBS†	TLB miss exception (Store)
4	AdEL	Address Error Exception (Load or instruction fetch)
5	AdES	Address Error Exception (Store)
6	IBE	Bus Error Exception (for Instruction Fetch)
7	DBE	Bus Error Exception (for data Load or Store)
8	Sys	SYSCALL Exception
9	Bp	Breakpoint Exception
10	RI	Reserved Instruction Exception
11	CpU	Co-Processor Unusable Exception
12	Ovf	Arithmetic Overflow Exception
13-31	-	Reserved

†These exceptions will not occur in R36100.

Table 5.12 Cause Register Exception Codes

The meanings of the other cause register bits are as follows:

- BD** The Branch Delay bit is set (1) if the last exception was taken while the processor was executing in the branch delay slot. If so, then the EPC will be rolled back to point to the branch instruction, so that it can be re-executed and the branch direction re-determined.
- CE** The Coprocessor Error field captures the coprocessor unit number referenced when a Coprocessor Unusable exception is detected.
- IP** The Interrupt Pending field indicates which interrupts are pending. Regardless of which interrupts are masked, the IP field can be used to determine which interrupts are pending.
- SW** The Software interrupt bits can be thought of as the logical extension of the IP field. The SW interrupts can be written to force an interrupt to be pending to the processor, and are useful in the prioritization of exceptions. To set a software interrupt, a "1" is written to the appropriate SW bit, and a "0" will clear the pending interrupt. There are corresponding interrupt mask bits in the status register for these interrupts.
- ExcCode** The exception code field indicates the reason for the last exception. Its values are listed in Table 5.12 on page 7.

The EPC (Exception Program Counter) Register

The 32-bit EPC register contains the virtual address of the instruction which took the exception, from which point processing resumes after the exception has been serviced. When the virtual address of the instruction resides in a branch delay slot, the EPC contains the virtual address of the instruction immediately preceding the exception (that is, the EPC points to the Branch or Jump instruction).

Bad VAddr Register

The Bad VAddr register saves the entire bad virtual address for any addressing exception.

The Status Register

The Status register contains all the major status bits; any exception puts the system in Kernel mode. All bits in the status register, with the exception of the TS (TLB Shutdown) bit, are readable and writable; the TS bit is read-only. Figure 5.4 on page 9 shows the functions of the various bits in the status register.

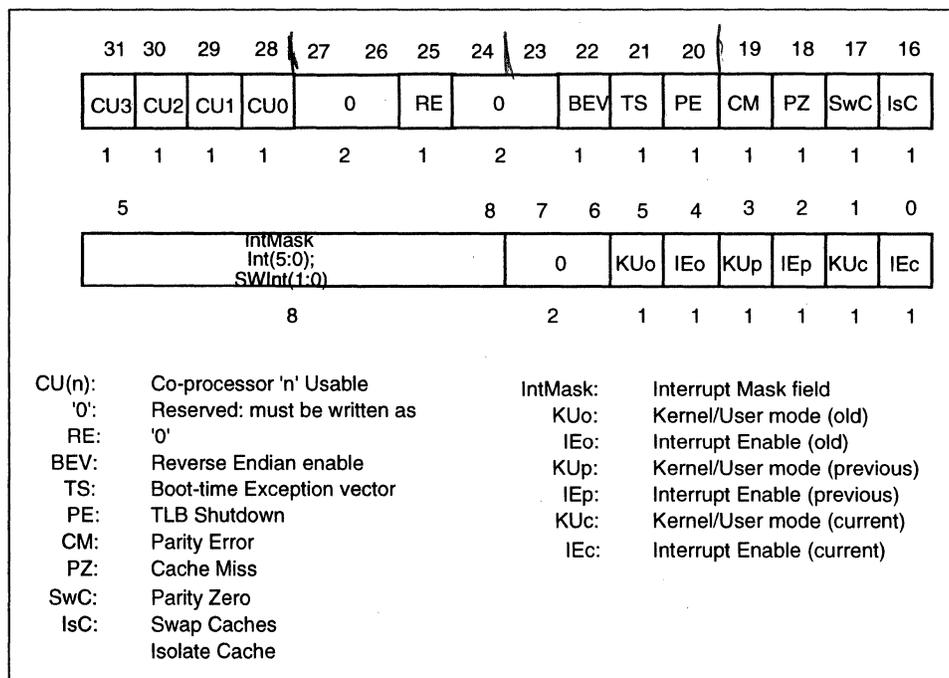


Figure 5.4 R36100 Status Register

The status register contains a three-level stack (current, previous, and old) of the kernel/user mode bit (KU) and the interrupt enable (IE) bit. The stack is pushed when each exception is taken and popped by the Restore From Exception instruction. These bits may also be directly read or written.

At reset, the SWc, KUc, and IEc bits are set to zero; BEV is set to one; and the value of the TS bit is set to "1". The rest of the bit fields are undefined after reset.

The various bits of the status register are defined as follows:

- CU** Coprocessor Usability. These bits individually control user level access to coprocessor operations, including the polling of the BrCond input pins and the manipulation of the System Control Coprocessor (CPO).
- RE** Reverse Endianness. The R3000 architecture allows the system to determine the byte ordering convention for the Kernel mode, and the default setting for user mode, at reset time. If this bit is cleared, the endianness defined at reset is used for the current user task. If this bit is set, then the user task will operate with the opposite byte ordering convention from that determined at reset. This bit has no effect on kernel mode. Also note that the setting of this bit does not affect the byte lanes used in 16- and 8-bit memory ports; thus, external byte lane shift logic is not required.
- BEV** Bootstrap Exception Vector. The value of this bit determines the locations of the exception vectors of the processor. If BEV = 1, then the processor is in "Bootstrap" mode, and the exception vectors reside in uncacheable space. If BEV = 0, then the processor is in normal mode, and the exception vectors reside in cacheable space.

TS	TLB Shutdown. This bit reflects whether the TLB is functioning. At reset, this bit can be used to determine whether the current processor is a base or extended architecture version. For the R36100, this bit is frozen at "1".
PE	Parity Error. This field should be written with a "1" at boot time. Once initialized, this field will always be read as "0".
CM	Cache Miss. This bit is set if a cache miss occurred while the cache was isolated. It is useful in determining the size and operation of the internal cache subsystem.
PZ	Parity Zero. This field should always be written with a "0".
SwC	Swap Caches. Setting this bit causes the execution core to use the on-chip instruction cache as a data cache and vice-versa. Resetting the bit to zero un-swaps the caches. This is useful for certain operations such as instruction cache flushing. This feature is not intended for normal operation with the caches swapped.
IsC	Isolate Cache. If this bit is set, the data cache is "isolated" from main memory; that is, store operations modify the data cache but do not cause a main memory write to occur, and load operations return the data value from the cache whether or not a cache hit occurred. This bit is also useful in various operations such as flushing, as described in Chapter 3.
IntMask	Interrupt Mask. This 8-bit field can be used to mask the hardware and software interrupts to the execution engine (that is, not allow them to cause an exception). IM(1:0) are used to mask the software interrupts, and IM (7:2) mask the 6 external interrupts. A value of '0' disables a particular interrupt, and a '1' enables it. Note that the IE bit is a global interrupt enable; that is, if the IE is used to disable interrupts, the value of particular mask bits is irrelevant; if IE enables interrupts, then a particular interrupt is selectively masked by this field.
KUo	Kernel/User old. This is the privilege state two exceptions previously. A '0' indicates kernel mode.
IEo	Interrupt Enable old. This is the global interrupt enable state two exceptions previously. A '1' indicates that interrupts were enabled, subject to the IM mask.
KUp	Kernel/User previous. This is the privilege state prior to the current exception. A '0' indicates kernel mode.
IEp	Interrupt Enable previous. This is the global interrupt enable state prior to the current exception. A '1' indicates that interrupts were enabled, subject to the IM mask.
KUc	Kernel/User current. This is the current privilege state. A '0' indicates kernel mode.
IEc	Interrupt Enable current. This is the current global interrupt enable state. A '1' indicates that interrupts are enabled, subject to the IM mask.
'0'	Fields indicated as '0' are reserved; they must be written as '0', and will return '0' when read.

PRID Register

This register is useful to software in determining which revision of the processor is executing the code. The format of this register is illustrated in Figure 5.5. For the R36100, the value returned is 0x0000_0710. On the R36100, the most significant 4 bits of the Revision field form an extension to the Implementation field. The least significant 4 bits of the Revision field are reserved for manufacturing. This value is different from other members of the IDT RISController family, so that software can easily determine the CPU type. This facilitates the development of one binary working with all family members.

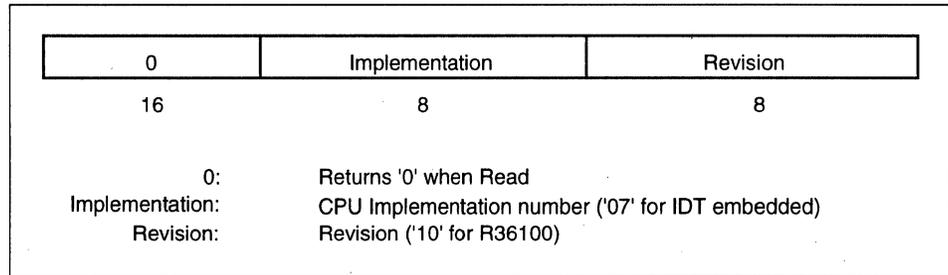


Figure 5.5 R36100 PrID Register



Integrated Device Technology, Inc.

Interrupt and Exception Handling

Chapter 6

Introduction

Processors execute code in a highly-directed fashion. The instruction immediately subsequent to the current instruction is fetched and then executed; if that instruction is a branch instruction, the program execution is diverted to the specified location. Thus, program execution is relatively straightforward and predictable.

Exceptions are a mechanism used to break into this execution stream and to force the processor to begin handling another task, typically related to either the system state or to the erroneous or undesirable execution of the program stream. Thus, exceptions typically are viewed by programmers as asynchronous interruptions of their program. (Note that exceptions are not necessarily unpredictable or asynchronous, in that the events which cause the exception may be exactly repeatable by the same software executing on the same data; however, the programmer does not typically "expect" an exception to occur when and where it does, and thus will view exceptions as asynchronous events).

The R3000 architecture provides for extremely fast, flexible interrupt and exception handling. The processor makes no assumptions about interrupt causes or handling techniques, and allows the system designer to build his own model of the best response to exception conditions. However, the processor provides enough information and resources to minimize both the amount of time required to begin handling the specific cause of the exception, and to minimize the amount of software required to preserve processor state information so that the normal instruction stream may be resumed.

This chapter discusses exception handling issues in R36100-based systems. The topics examined are: the exception model, the machine state to be saved on an exception, and nested exceptions. Representative software examples of exception handlers are also provided, as are techniques and issues appropriate to specific classes of exceptions.

R36100 Exception Model

The exception processing capability of the R36100 assures an orderly transfer of control from an executing program to the kernel. Exceptions may be broadly divided into two categories: they can be caused by an instruction or instruction sequence, including an unusual condition arising during its execution; or can be caused by external events such as interrupts. When an R36100 detects an exception, the normal sequence of instruction flow is suspended; the processor is forced to kernel mode where it can respond to the abnormal or asynchronous event. Table 6.1 on page 2 lists the exceptions recognized by the R3000 architecture.

Exception	Mnemonic	Cause
Reset	Reset	Assertion of the Reset signal causes an exception that transfers control to the special vector at virtual address 0xbfc0_0000.
UTLB Miss†	UTLB	User TLB Miss. A reference is made (in either kernel or user mode) to a page in kuseg that has no matching TLB entry. This can occur only in extended architecture versions of the processor.
TLB Miss†	TLBL (Load) TLBS (Store)	A referenced TLB entry's Valid bit isn't set, or there is a reference to a kseg2 page that has no matching TLB entry. This can occur only in extended architecture versions of the processor.
TLB Modified†	Mod	During a store instruction, the Valid bit is set but the dirty bit is not set in a matching TLB entry. This can occur only in extended architecture versions of the processor.
Bus Error	IBE (Instruction) DBE (Data)	Assertion of the Bus Error input during a read operation, due to such external events as bus timeout, backplane memory errors, invalid physical address, or invalid access types.
Address Error	AdEL (Load) AdES (Store)	Attempt to load, fetch, or store an unaligned word; that is, a word or halfword at an address not evenly divisible by four or two, respectively. Also caused by reference to a virtual address with most significant bit set while in User Mode.
Overflow	Ovf	Twos complement overflow during add or subtract.
System Call	Sys	Execution of the SYSCALL Trap Instruction
Breakpoint	Bp	Execution of the break instruction
Reserved Instruction	Rl	Execution of an instruction with an undefined or reserved major operation code (bits 31:26), or a special instruction whose minor opcode (bits 5:0) is undefined.
Co-processor Unusable	CpU	Execution of a co-processor instruction when the CU (Co-processor Usable) bit is not set for the target co-processor.
Interrupt	Int	Assertion of one of the six hardware interrupt inputs or setting of one of the two software interrupt bits in the Cause register.
†These exceptions will not occur in an R36100, or in any base member of the R30xx family.		

Table 6.1 R3000 Architecture Exceptions

Precise vs. Imprecise Exceptions

One classification of exceptions refers to the precision with which the exception cause and processor context can be determined. That is, some exceptions are precise in their nature, while others are "imprecise."

In a **precise exception**, much is known about the system state at the exact instance the exception is caused. Specifically, the exact processor context and the exact cause of the exception are known. The processor thus maintains its exact state before the exception was generated, and can accurately handle the exception, allowing the instruction stream to resume when the situation is corrected. Additionally, in a precise exception model, the processor can not advance state; that is, subsequent instructions, which may already be in the processor pipeline, are not allowed to change the state of the machine.

Many real-time applications greatly benefit from a processor model which guarantees precise exception context and cause information. The MIPS architecture, including the R36100, implements a precise exception model for all exceptional events.

Exception Processing

The R36100 exception handling system efficiently handles machine exceptions, including arithmetic overflows, I/O interrupts, system calls, breakpoints, reset, and co-processor unusable conditions. Any of these events interrupt the normal execution flow; the R36100 aborts the instruction causing the exception and also aborts all those following in the exception pipeline which have already begun, thus not modifying processor context. The CPU then performs a direct jump into a designated exception handler routine. This insures that the R36100 is always consistent with the precise exception model.

Exception Handling Registers

The system co-processor (CPO) registers contain information pertinent to exception processing. Software can examine these registers during exception processing to determine the cause of the exception and the state of the processor when it occurred. There are four registers used in exception processing, shown in Chapter 5. These are the *Cause* register, the *EPC* register, the *Status* register, and the *BadVAddr* register. A brief description of each follows.

The Cause Register

The contents of the Cause register describe the last exception. A 5-bit exception code indicates the cause of the current exception; the remaining fields contain detailed information specific to certain exceptions.

All bits in this register, with the exception of the SW bits, are read-only. The SW bits can be written to set or reset software interrupts. Figure 6.1 shows the cause register.

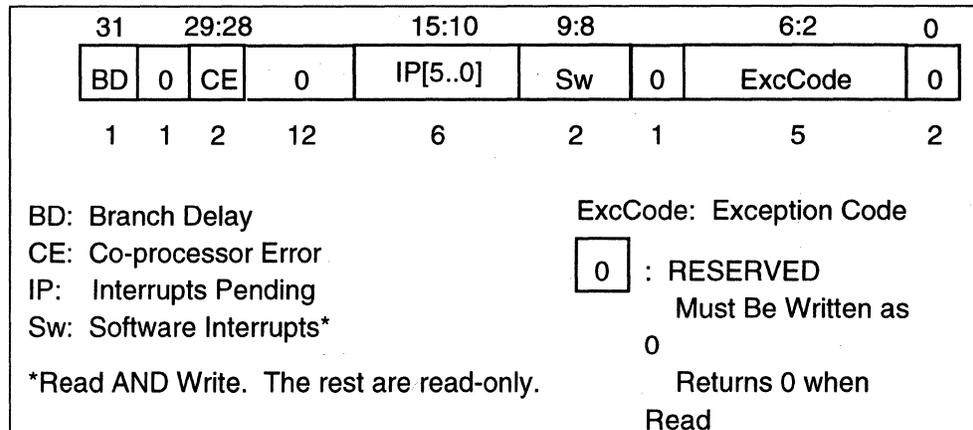


Figure 6.1 R36100 Cause Register

The meaning of the other bits of the cause register is as follows:

- BD** The Branch Delay bit is set (1) if the last exception was taken while the processor was executing in the branch delay slot. If so, then the EPC will be rolled back to point to the branch instruction, so that it can be re-executed and the branch direction re-determined.
- CE** The Co-processor Error field captures the co-processor unit number referenced when a Co-processor Unusable exception is detected.
- IP** The Interrupt Pending field indicates which interrupts are pending. Regardless of which interrupts are masked, the IP field can be used to determine which interrupts are pending.
- SW** The Software interrupt bits can be thought of as the logical extension of the IP field. The SW interrupts can be written to force an interrupt to be pending to the processor, and are useful in the prioritization of exceptions. To set a software interrupt, a "1" is written to the appropriate SW bit, and a "0" will clear the pending interrupt. There are corresponding interrupt mask bits in the status register for these interrupts.
- ExcCode** The exception code field indicates the reason for the last exception. Its values are listed in Table 6.2.

Number	Mnemonic	Description
0	Int	External Interrupt
1	MOD†	TLB Modification Exception
2	TLBL†	TLB miss Exception (Load or instruction fetch)
3	TLBS†	TLB miss exception (Store)
4	AdEL	Address Error Exception (Load or instruction fetch)
5	AdES	Address Error Exception (Store)
6	IBE	Bus Error Exception (for Instruction Fetch)
7	DBE	Bus Error Exception (for data Load or Store)
8	Sys	SYSCALL Exception
9	Bp	Breakpoint Exception
10	RI	Reserved Instruction Exception
11	CpU	Co-Processor Unusable Exception
12	Ovf	Arithmetic Overflow Exception
13-31	-	Reserved
†These exceptions will not occur in an R36100		

Table 6.2 Cause Register Exception Codes

The EPC (Exception Program Counter) Register

The 32-bit EPC register contains the virtual address of the instruction which took the exception, from which point processing resumes after the exception has been serviced. When the virtual address of the instruction resides in a branch delay slot, the EPC contains the virtual address of the instruction immediately preceding the exception (that is, the EPC points to the Branch or Jump instruction).

Bad VAddr Register

The Bad VAddr register saves the entire bad virtual address for any addressing exception.

The Status Register

The Status register contains all the major status bits; any exception puts the system in Kernel mode. All bits in the status register, with the exception of the TS (TLB Shutdown) bit, are readable and writable; the TS bit is read-only, and frozen to '1' in the R36100. Figure 6.2 shows the definition and position of the various bits in the status register.

The status register contains a three level stack (current, previous, and old) of the kernel/user mode bit (KU) and the interrupt enable (IE) bit. The stack is pushed when each exception is taken, and popped by the Restore From Exception instruction. These bits may also be directly read or written.

At reset, the SWc, KUc, and IEc bits are set to zero; BEV is set to one; and the value of the TS bit is set to "1". The rest of the bit fields are undefined after reset.

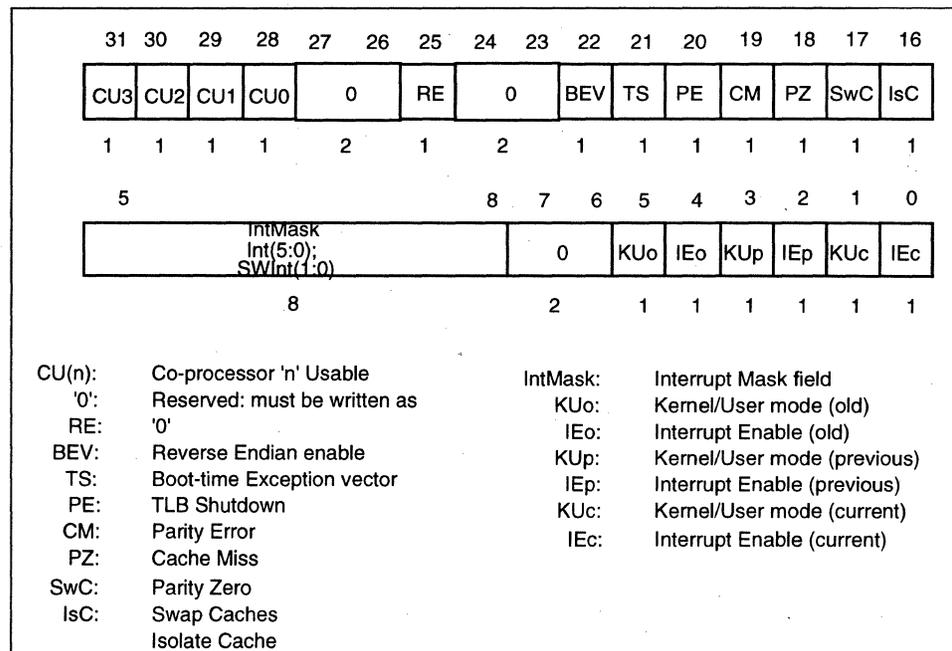


Figure 6.2 The Status Register

The various bits of the status register are defined in chapter 5. The bits of most relevance in exception processing are repeated below.

- BEV** Bootstrap Exception Vector. The value of this bit determines the locations of the exception vectors of the processor. If BEV = 1, then the processor is in "Bootstrap" mode, and the exception vectors reside in uncacheable space. If BEV = 0, then the processor is in normal mode, and the exception vectors reside in cacheable space.
- IM** Interrupt Mask. This 8-bit field can be used to mask the hardware and software interrupts to the execution engine (that is, not allow them to cause an exception). IM(1:0) are used to mask the software interrupts, and IM (7:2) mask the 6 external interrupts. A value of '0' disables a particular interrupt, and a '1' enables it. Note that the IE bit is a global interrupt enable; that is, if the IE is used to disable interrupts, the value of particular mask bits is irrelevant; if IE enables interrupts, then a particular interrupt is selectively masked by this field.
- KUo** Kernel/User old. This is the privilege state two exceptions previously. A '0' indicates kernel mode.
- IEo** Interrupt Enable old. This is the global interrupt enable state two exceptions previously. A '1' indicates that interrupts were enabled, subject to the IM mask.
- KUp** Kernel/User previous. This is the privilege state prior to the current exception. A '0' indicates kernel mode.
- IEp** Interrupt Enable previous. This is the global interrupt enable state prior to the current exception. A '1' indicates that interrupts were enabled, subject to the IM mask.

- KUc** Kernel/User current. This is the current privilege state. A '0' indicates kernel mode.
- IEc** Interrupt Enable current. This is the current global interrupt enable state. A '1' indicates that interrupts are enabled, subject to the IM mask.

Exception Vector Locations

The R3000 architecture separates exceptions into three vector spaces. The value of each vector depends on the BEV (Boot Exception Vector) bit of the status register, which allows two alternate sets of vectors (and thus two different pieces of code) to be used.

Typically, this is used to allow diagnostic tests to occur before the functionality of the cache is validated; processor reset forces the value of the BEV bit to a '1'. Table 6.3 and Table 6.4 list the exception vectors for the R36100 for the two different modes.

Exception	Virtual Address	Physical Address
Reset	0xbfc0_0000	0x1fc0_0000
UTLB Miss	0x8000_0000	0x0000_0000
General	0x8000_0080	0x0000_0080

Table 6.3 Exception Vectors When BEV = 0

Exception	Virtual Address	Physical Address
Reset	0xbfc0_0000	0x1fc0_0000
UTLB Miss	0xbfc0_0100	0x1fc0_0100
General	0xbfc0_0180	0x1fc0_0180

Table 6.4 Exception Vectors When BEV = 1

Exception Prioritization

It is important to understand the structure of the R36100 instruction execution unit in order to understand the exception priority model of the processor. The R36100 runs instructions through a five stage pipeline, illustrated in Figure 6.3.

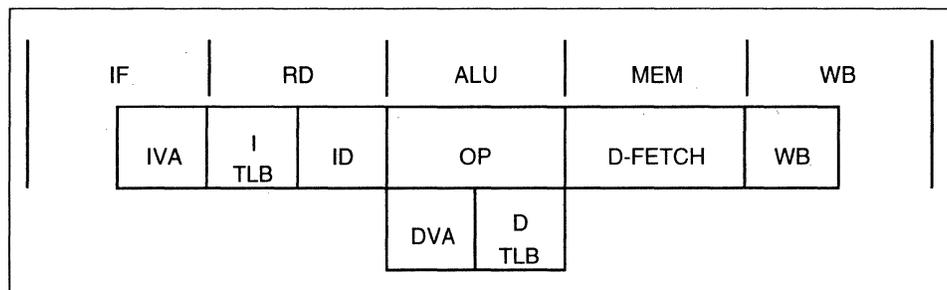


Figure 6.3 Pipelining in the R3051 family

The pipeline stages are as follows:

- IF (Instruction Fetch). This cycle contains two parts: the IVA (Instruction Virtual Address) phase, which generates the virtual instruction address of the next instruction to be fetched, and the ITLB phase, which performs the virtual to physical translation of the address.
- RD (Read and Decode). This phase obtains the required data from the internal registers and also decodes the instruction.
- ALU (Arithmetic Logic Unit). This phase either performs the desired arithmetic or logical operation, or generates the address for the upcoming data operation. For data operations, this phase contains both the data virtual address stage, which generates the desired virtual address, and the data TLB stage, which performs the virtual to physical translation.
- MEM (Memory). This phase performs the data load or store transaction.
- WB (Write Back). This stage updates the registers with the result data.

High performance is achieved because five instructions are operating concurrently, each in a different stage of the pipeline. However, since multiple instructions are operating concurrently, it is possible that multiple exceptions are generated concurrently. If so, the processor must decide which exception to process, basing this decision on the stage of the pipeline that detected the exception. The processor will then flush all preceding pipeline stages to avoid altering processor context, thus implementing precise exceptions. This determines the relative priority of the exceptions.

For example, an illegal instruction exception can only be detected in the instruction decode stage of the R36100; an Instruction Bus Error can only be determined in the I-Fetch pipe stage. Since the illegal instruction was fetched before the instruction which generated the bus error was fetched, and since it is conceivable that handling this exception might have avoided the second exception, it is important that the processor handle the illegal instruction before the bus error. Therefore the exception detected in the latest pipeline stage has priority over exceptions detected in earlier pipeline stages. All instructions fetched subsequent to this (all preceding pipeline stages) are flushed to avoid altering state information, maintaining the precise exception model.

Table 6.5 lists the priority of exceptions from highest first to lowest.

Mnemonic	Pipestage
Reset	Any
AdEL	Memory (Load instruction)
AdES	Memory (Store instruction)
DBE	Memory (Load or store)
MOD†	ALU (Data TLB)
TLBL†	ALU (DTLB Miss)
TLBS†	ALU (DTLB Miss)
Ovf	ALU
Int	ALU
Sys	RD (Instruction Decode)
Bp	RD (Instruction Decode)
RI	RD (Instruction Decode)
CpU	RD (Instruction Decode)
TLBL†	I-Fetch (ITLB Miss)
AdEL	IVA (Instruction Virtual Address)
IBE	RD (end of I-Fetch)
†These exceptions will not occur in an R36100, which does not include a TLB.	

Table 6.5 R36100 Exception Priority

Exception Latency

A critical measurement of a processor’s throughput in interrupt driven systems is the interrupt “latency” of the system. Interrupt latency is a measurement of the amount time from the assertion of an interrupt until software begins handling that interrupt. Often included when discussing latency is the amount of overhead associated with restoring context once the exception is handled, although this is typically less critical than the initial latency.

In systems where the processor is responsible for managing a number of time-critical operations in real time, it is important that the processor minimize interrupt latency. That is, it is more important that *every* interrupt be handled at a rate above some given value, rather than *occasionally* handle an interrupt at very high speed.

Factors which affect the interrupt latency of a system include the types of operations it performs (that is, systems which have long sequences of operations during which interrupts can not be accepted have long latency), how much information must be stored and restored to preserve and restore processor context, and the priority scheme of the system.

Table 6.5 illustrates which pipestage recognizes which exceptions. As mentioned above, all instructions less advanced in the pipeline are flushed from the pipeline to avoid altering state execution. Those instructions will be restarted when the exception handler completes.

Once the exception is recognized, the address of the appropriate exception vector will be the next instruction to be fetched. In general, the latency to the exception handler is one instruction cycle, and at worst the longest stall cycle in that system.

The R36100 implements mechanisms which can help improve exception response time. Primary among these is the cache locking mechanism described in earlier chapters. System software can be easily arranged such that the exception service routines and/or critical exception data are locked into the on-chip cache. The result will be both high-speed and fully deterministic.

Interrupts Inputs in the R36100

The organization of interrupts in an R36100-based system is up to the system architect. Specifically, the R36100 multiplexes various interrupt pins with PIO pins; depending on the programming of the PIO unit, the system may have 6 external interrupts and 2 BrCond input pins available for interrupt software. This section describes operation assuming all such inputs are available to system software. Later chapters describe the on-chip PIO and interrupt control units.

Interrupt Operation in the R36100

The R36100 family features two types of interrupt inputs: synchronized internally and non-synchronized, or direct.

The $\overline{SInt}(2:0)$ bus (Synchronized Interrupts) allow the system designer to connect unsynchronized interrupt sources to the processor. The processor includes special logic on these inputs to avoid meta-stable states associated with switching inputs right at the processor sampling point. Because of this logic, these interrupt sources have slightly longer latency from the $\overline{SInt}(n)$ pin to the exception vector than the non-synchronized inputs. The operation of the synchronized interrupts is illustrated in Figure 6.4.

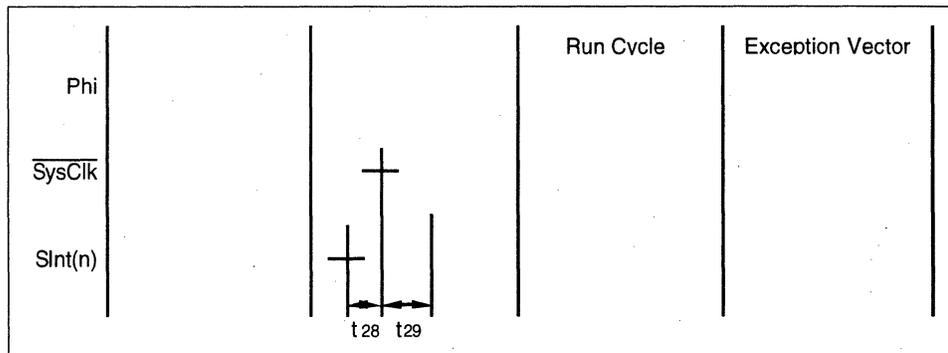


Figure 6.4 Synchronized Interrupt Operation

The other interrupts, $\overline{Int}(5:3)$, do not contain this synchronization logic, and thus have slightly better latency to the exception vector. However, the interrupting agent must guarantee that it always meets the interrupt input set-up and hold time requirements of the processor. These inputs are useful for interrupting agents which operate off of the synchronously with the R36100. The operation of these interrupts is illustrated in Figure 6.5.

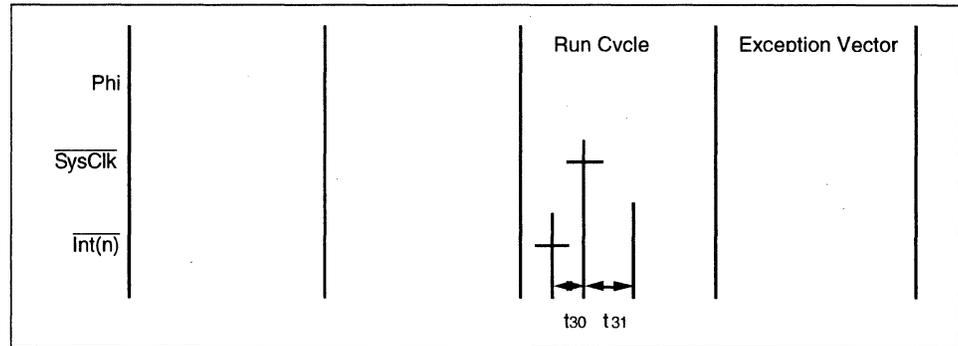


Figure 6.5 Direct Interrupt Operation

Since the interrupt exception is detected during the ALU stage of the instruction currently in the processor pipeline, at least one run cycle must occur between (or at) the assertion of the external interrupt input and the fetch of the exception vector. Thus, if the processor is in a stall cycle when an external agent sends an interrupt, it will execute at least one run cycle before beginning exception processing. In this instance, there would be no difference in the latency of synchronized and direct interrupt inputs.

All of the interrupts are level-sensitive and active low. They continue to be sampled after an interrupt exception has occurred, and are not latched within the processor when an interrupt exception occurs. It is important that the external interrupting agent maintain the interrupt line until software acknowledges the interrupt.

Note that the R3081 incorporates a hardware floating point accelerator on-chip. The MIPS architecture recommends that Int(3) be used to handle the floating point interrupt; thus, the R3081 defaults to this interrupt assignment. However, the R3081 Config register (which differs from the R36100 Config register) can be used to change the assignment.

Also, the on-chip interrupt controller of the R36100 will signal its interrupt to the CPU using one of the available CPU interrupts. The interrupt controller defaults to Int(4) for this operation; however, the interrupt controller does allow software to select an alternative interrupt. In any case, the system needs to reserve one CPU interrupt for the on-chip interrupt controller.

Each of the eight interrupts (6 hardware and 2 software) can be individually masked by clearing the corresponding bit in the Interrupt Mask field of the Status Register. All eight interrupts can be masked at once by clearing the IEC bit in the Status Register.

On the synchronized interrupts, care should be taken to allow at least two clock cycles between the negation of the interrupt input and the re-enabling of the interrupt mask for that bit. In general, it is recommended that software continue polling the IP field of the Cause register once it has instructed the peripheral to negate its interrupt, prior to re-enabling its mask, to avoid a spurious interrupt.

The value shown in the interrupt pending bits of the Cause register reflects the current state of the interrupt pins of the processor. These bits are not latched (except for sampling from the data bus to guarantee that they are stable when examined), and the masking of specific interrupt inputs does not mask the bits from being read.

Using the BrCond Inputs

In addition to the interrupt pins themselves, many systems can use the BrCond(3:2) input port pins in their exception model. These pins can be directly tested by software, and can be used for polling or fast interrupt decoding. The kernel must enable the use of the corresponding co-processor unit before testing the state of the BrCond input pin.

The R36100 provides two synchronized BrCond inputs: SBrCond(3:2). Note that BrCond(0), corresponding to the on-chip CP0, and BrCond(1), corresponding to Co-Processor 1 (the FPA, present on the R3081), are not available on the R36100 as user inputs. Instructions that use BrCond(1:0) will always see a '1' on the R36100. Also note that the SBrCond(3:2) on the R36100 may be not be enabled in the PIO unit, in which case the SBrCond(3:2) input values are undefined. When programmed to be SBrCond(3:2) inputs, the timing requirements of the SBrCond inputs are illustrated in Figure 6.6. Since these inputs are synchronized by the R36100, they do not need to be driven synchronously to the processor.

Similar to the interrupt inputs, at least one instruction must be executed (in the ALU stage) of the instruction pipeline prior to software being able to detect a change in one of these inputs. This is because the processor actually captures the value of these flags one instruction prior to the branch on co-processor instruction. Before executing a Branch Condition instruction (i.e. BCzT, BCzF) the corresponding co-processor usable bit in the CPO status register must be set; otherwise, a co-processor unusable exception will be signalled.

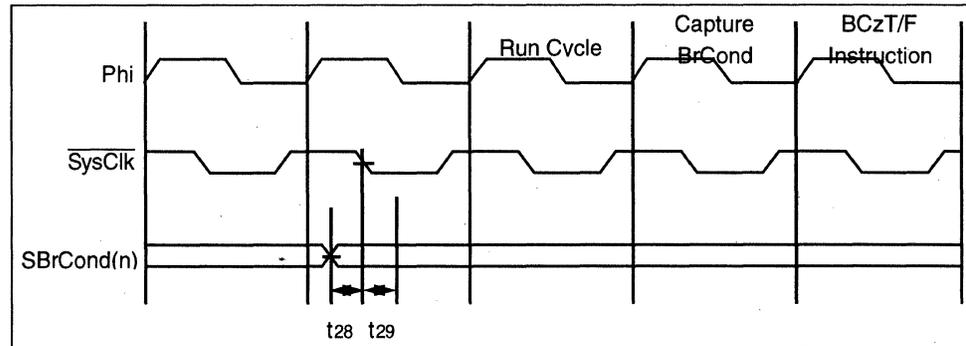


Figure 6.6 Synchronized BrCond Inputs

Interrupt Handling

The assertion of an unmasked interrupt input causes the R36100 to branch to the general exception vector at virtual address 0x8000_0080, and write the 'Int' code in the Cause register. The IP field of the Cause register shows which of the six hardware interrupts are pending and the SW field in the Cause register show which of the two software interrupts are pending. Multiple interrupts can be pending at the same time, with no priority assumed by the processor.

If the interrupt asserted is due to the on-chip interrupt controller, the interrupt controller must be accessed to determine which of its interrupt sources caused the assertion. This operation is described in a later chapter.

When an interrupt occurs, the KUp, IEp, KUc and IEc bits of the Status register are saved in the KUo, IEo, KUp, IEp bit fields in the Status register, respectively, as illustrated in Figure 6.7. The current kernel status bit KUc and the interrupt bit IEc are cleared. This will mask all of the interrupts and then place the processor in kernel mode. This sequence will be reversed by the execution of an *rfe* (restore from exception) instruction, typically in the branch delay slot of the branch which resumes normal execution.

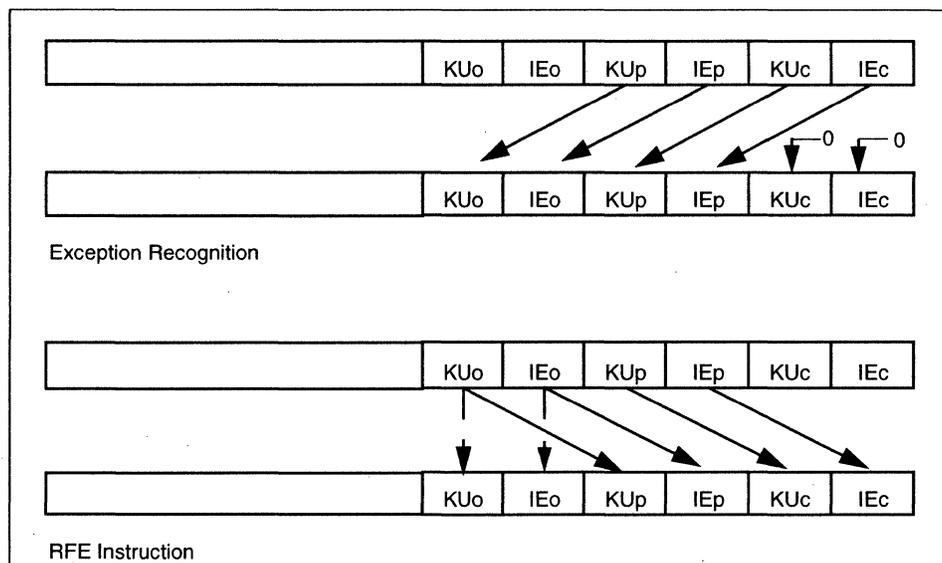


Figure 6.7 Kernel and Interrupt Status Being Saved on Interrupts

Interrupt Servicing

In case of a hardware interrupt, the interrupt must be cleared by de-asserting the interrupt line, which has to be done by alleviating the external conditions that caused the interrupt. Software interrupts have to be cleared by clearing the corresponding bits, SW(1:0), in the Cause register to zero. It is recommended that software continue polling the IP field of the Cause register once it has instructed the peripheral to negate its interrupt, prior to re-enabling its mask, to avoid a spurious interrupt.

Basic Software Techniques For Handling Interrupts

Once an exception is detected the processor suspends the current task, enters kernel mode, disables interrupts, and begins processing at the exception vector location. The EPC is loaded with the address the processor will return to once the exception event is handled.

The specific actions of the processor depend on the cause of the exception being handled. The MIPS architecture classifies exceptions into three distinct classes: RESET, UTLB Miss, and General.

Coming out of reset, the processor initializes the state of the machine. In addition to initializing system peripherals, page tables, the TLB, and the caches, software clears both STATUS and CAUSE registers, and initializes the exception vectors.

The code located at the exception vector may be just a branch to the actual exception code; however, in more time critical systems the instructions located at the exception vector may perform the actual exception processing. In order to cause the exception vector location to branch to the appropriate exception handler (presuming that such a jump is appropriate), a short code sequence such as that illustrated in Figure 6.8 may be used.

It should be noted the contents of register k0 are not preserved. This is not a problem for software, since MIPS compiler and assembler conventions reserve k0 (and often k1) for kernel processes, and do not use it for user programs. For the system developer it is advised that the use of k0 be reserved for use by the exception handling code exclusively. This will make debugging and development much easier.

The "IDT R30xx Family Software Reference Manual" provides a great deal of information on the software requirements of exception management, including interrupt service.

```

.set    noreorder    # tells the assembler not to reorder the code
/*
**   code sequence copied to UTLB exception vector
*/
    la    k0,excep_utlb    #address of utlb excp. handler
    j    k0                # jump via reg k0
    nop
/*
**   code sequence copied to general exception vector
*/
    la    k0,excep_gener    #address of general excp. handler
    al
    j    k0                # jump via reg k0
    nop

```

Figure 6.8 Code Sequence to Initialize Exception Vectors

Preserving Context

The R36100 has the following four registers related to exception processing:

- The *Cause* register
- The *EPC* (exception program counter) register
- The *Status* register
- The *BadVAddr* (bad virtual address) register

Typical exception handlers preserve the status, cause, and EPC registers in general registers (or on the system stack). If the exception cause is due to an address error, software may also preserve the bad virtual address register for later processing.

Note that not all systems need to preserve this information. Since the R36100 disables subsequent interrupts, it is possible for software to directly process the exception while leaving the processor context in the CPO registers. Care must be taken to insure that the execution of the exception handler does not generate subsequent exceptions.

Preserving the context in general registers (and on the stack) does have the advantage that interrupts can be re-enabled while the original exception is handled, thus allowing a priority interrupt model to be built.

A typical code sequence to preserve processor context is shown in Figure 6.9. This code sequence preserves the context into an area of memory pointed to by the k0 kernel register. This register points to a block of memory capable of storing processor context. Constants identified by name (such as *R_EPC*) are used to indicate the offset of a particular register from the start of that memory area.

It should be noted that this sequence for fetching the co-processor zero registers is required because there is a one clock delay in the register value actually being loaded into the general registers after the execution of the *mfc0* instruction.

```

la      k0,except_regs    # fetch address of reg save array
sw      AT,R_AT*4(k0)     # save register AT
sw      v0,R_V0*4(k0)     # save register v0
sw      v1,R_V1*4(k0)     # save register v1
mfc0    v0,C0_EPC         # fetch the epc register
mfc0    v1,C0_SR          # fetch the status register
sw      v0,R_EPC*4(k0)    # save the epc
mfc0    v0,C0_CAUSE       # fetch the cause register
sw      v1,R_SR*4(k0)     # save status register

/*      The above code is about the minimum required
**      The user specific code would follow
*/

```

Figure 6.9 Preserving Processor Context

Determining The Cause Of The Exception

The cause register indicates the reason the exception handler was invoked. Thus, to invoke the appropriate exception service routine, software merely needs to examine the cause register, and use its contents to direct a branch to the appropriate handler.

One method of decoding the jump to an appropriate software routine to handle the exception and cause is shown in Figure 6.10. Register `v0` contains the cause register, and register `k0` still points to the register save array.

<code>.set</code>	<code>noreorder</code>	
<code>sw</code>	<code>a0,R_A0*4(k0)</code>	<code># save register a0</code>
<code>and</code>	<code>v1,v0,EXCMASK</code>	<code># isolate exception code</code>
<code>lw</code>	<code>a0,cause_table(v1)</code>	<code># get address of interrupt routine.</code>
<code>sw</code>	<code>a1,R_A1*4(k0)</code>	<code># use delay slot to save register a1</code>
<code>j</code>	<code>a0</code>	
<code>sw</code>	<code>k1,R_K1*4(sp)</code>	<code># save k1 register</code>
<code>.set</code>	<code>reorder</code>	<code># re-enable pipeline scheduling</code>

Figure 6.10 Exception Cause Decoding

The above sequence of instructions extracts the exception code from the cause register and uses that code to index into the table of pointers to functions (the `cause_table`). The `cause_table` data structure is shown in Figure 6.11.

Each of the entries in this table point to a function for processing the particular type of interrupt detected. The specifics of the code contained in each function is unique to a given application; all registers used in these functions must be saved and restored.

```
int (*cause_table[16])() = {
    int_extern,      /* External interrupts      */
    int_tlbmod,     /* TLB modification error   */
    int_tlbmiss,    /* load or instruction fetch */
    int_tlbmiss,    /* write miss                */
    int_addrerr,    /* load or instruction fetch */
    int_addrerr,    /* write address error       */
    int_ibe,        /* Bus error - Instruction fetch */
    int_dbe,        /* Bus error - load or store data */
    int_syscall,    /* SYSCALL exception        */
    int_breakpoint, /* breakpoint instruction    */
    int_trap,       /* Reserved instruction      */
    int_cpunuse,    /* coprocessor unusable     */
    int_trap,       /* Arithmetic overflow       */
    int_unexp,      /* Reserved                   */
    int_unexp,      /* Reserved                   */
    int_unexp,      /* Reserved                   */
};
```

Figure 6.11 Exception Service Branch Table

Returning From Exceptions

Returning from the exception routine is made through the *rfe* instruction. When the exception first occurs the R36100 automatically saves some of the processor context, the current value of the interrupt enable bit is saved into the field for the previous interrupt enable bit, and the kernel/user mode context is preserved.

The *IE* interrupt enable bit must be asserted (a one) for external interrupts to be recognized. The *KU* kernel mode bit must be a zero in kernel mode. When an exception occurs, external interrupts are disabled and the processor is forced into kernel mode. When the *rfe* instruction is executed at completion of exception handling, the state of the mode bits is restored to what it was when the exception was recognized (presuming the programmer restored the status register to its value when the exception occurred). This is done by “popping” the old/previous/current *KU* and *IE* bits of the status register.

The code sequence in Figure 6.12 is an example of exiting an interrupt handler. The assumption is that registers and context were saved as outlined above. To properly exit from exception handling, this code sequence must either be replicated in each of the cause handling functions or each of them must branch to this code sequence.

Note that this code sequence must be executed with interrupts disabled. If the exception handler routine re-enables interrupts, they must be disabled when the *CPO* registers are being restored.

```

gen_excp_exit:
    .set    noreorder
                                     # by the time we have gotten here
                                     # all general registers have been
                                     # restored (except of k0 and v0)
                                     # reg. AT points to the reg save array
    lw     k0,CO_SR*4(AT)           # fetch status reg. contents
    lw     v0,R_V0*4(AT)           # restore reg. v0
    mtc0   k0,CO_SR                # restore the status reg. contents
    lw     k0,R_EPC*4(AT)          # Get the return address
    lw     AT,R_AT*4(AT)           # restore AT in load delay
    j      k0                       # return from int. via jump reg.
    rfe                                     # the rfe instr. is executed in the
                                     # branch delay slot
    .set    reorder

```

Figure 6.12 Returning from Exception

Special Techniques For Interrupt Handling

There are a number of techniques which take advantage of the R36100 architecture to minimize exception latency and maximize throughput in interrupt driven systems. This section discusses a number of those techniques.

Interrupt Masking

Only the six external and two software interrupts are maskable exceptions. The mask for these interrupts are in the status register.

To enable a given external interrupt, the corresponding bit in the status register must be set. The IEC bit in the status register must also be set. It follows that by setting and clearing these bits within the interrupt handler that interrupt priorities can be established. The general mechanism for doing this is performed within the external interrupt-handler portion of the exception handler.

The interrupt handler preserves the current mask value when the status register is preserved. The interrupt handler then calculates which (if any) external interrupts have priority, and sets the interrupt mask bit field of the status register accordingly. Once this is done, the IEC bit is changed to allow higher priority interrupts. Note that all interrupts must again be disabled when the return from exception is processed.

Using BrCond For Fast Response

The R36100 instruction set contains mechanisms to allow external or internal co-processors to operate as an extension of the main CPU. Some of these features may also be used in an interrupt-driven system to provide the highest levels of response.

Specifically, the R36100 allows external input port signals, the SBrCond(3:2) signals. These signals are used by external agents to report status back to the processor. The instruction set contains instructions which allow the external bits to be tested, and branches to be executed depending on the value of the SBrCond input.

An interrupt-driven system can use these SBrCond signals, and the corresponding instructions, to implement an input port for time-critical interrupts. Rather than mapping an input port in memory (which requires external logic), the SBrCond signals can be examined by software to control interrupt handling.

There are actually two techniques to use this advantageously. One method uses these signals to perform interrupt polling; in this method, the processor continually examines these signals, waiting for an appropriate value before handling the interrupt. A sample code sequence is shown in Figure 6.13.

The software in this system is very compact, and easily resides in the on-chip cache of the processor. Thus, the latency to the interrupt service routine in this system is minimized, allowing the fastest interrupt service capabilities.

A second method utilizes external interrupts combined with the SBrCond signals. In this method, both the SBrCond signal and one of the external interrupt lines are asserted when an external event occurs. This configuration allows the CPU to perform normal tasks while waiting for the external event.

For example, assume that a valve must be closed and then normal processing continued when SBrCond(2) is asserted TRUE. The valve is controlled by a register that is memory-mapped to address 0xaffe_0020 and writing a one to this location closes the valve. The software in Figure 6.14 accomplishes this, using SBrCond(2) to aid in cause decoding.

The number of cycles for a deterministic system is five cycles between the time the interrupt occurred and it was serviced. Interrupts were re-enabled in four additional cycles. Note that none of the processor context needs to be preserved and restored for this routine.

```
.set    noreorder    # prevents the assembler from
                    # reordering the code below

polling_loop:      # branch to yourself until
    bc2f    polling_loop    # BrCond(2) is asserted
    nop

                    # Once BrCond(2) is asserted, fall through
                    # and begin processing the external event

fast_response_cp2:

                    # code sequence that would do the
                    # event processing

    b      polling_loop    # return to polling
```

Figure 6.13 Polling System Using BrCond

```

.set    noreorder          # prevents the assembler from reordering
                               # the code sequences below

/* This section of code is placed at the general exception
** vector location 0x8000_0080. When an external interrupt is
** asserted execution begins here.

*/

        bc2t   close_valve   # test for emergency condition and
        li     k0,1          # jump to close valve if TRUE
        la     k0,gen_exp_hand # otherwise,
        j      k0            # jump to general exc. handler
        nop                               # and process less critical excepts.

/* This is the close valve routine - its sole purpose is to close the
** valve as quickly as possible. The registers 'k0' and 'k1' are reserved
** for kernel use and therefore need not be saved when a client or
** user program is interrupted. It should be noted that the value to
** write to the valve close register was put in reg 'k0' in the
** branch delay slot above - so by the time we get here it is
** ready to output to the close register.

*/

close_valve:
        la     k1,0xaffe0020 # the address of the close register
        sw    k0,0(k1)       # write the value to the close register
        mfc0  k0,CO_EPC      # get the return address to cont processing
        nop
        j     k0             # return to normal processing
        rfe                               # restore previous interrupt mask
                               # and kernel/user mode bits of the
                               # status register.

.set    reorder

```

Figure 6.14 Using BrCond for Fast Interrupt Decoding

Cache Locking

The R36100 allows the cache to be split into multiple sections, each servicing a different section of the processor address space. Using this technique, the system can "dedicate" one such section to exception service. Since the portion of the cache which deals with exception service will not be disturbed by normal system operation, the exception service code can be effectively locked into the on-chip cache. This has two positive benefits.

First, this insures that the exception service routine will operate directly out of the on-chip cache, and avoid main memory I-cache miss fetches. This speeds overall execution.

Secondly, this insures that the exception service routine performance will not be dependent on the tasks run since it was last invoked. Since those tasks will not displace the exception software from the on-chip cache, the exception software performance will be deterministic.

Nested Interrupts

Note that the processor does not automatically stack processor context when an exception occurs; thus, to allow nested exceptions it is important that software perform this stacking.

Most of the software illustrated above also applies to a nested exception system. However, rather than using just one register (pointed to by k0) as a save area, a stacking area must be implemented and managed by software. Also, since interrupts are automatically disabled once an exception is detected, the interrupt handling routine must mask the interrupt it is currently servicing, and re-enable other interrupts (once context is preserved) through the IEC bit.

The use of Interrupt Mask bits of the status register to implement an interrupt prioritization scheme was discussed earlier. An analogous technique can be performed by using an external interrupt encoder to allow more interrupt sources to be presented to the processor.

Software interrupts can also be used as part of the prioritization of interrupts. If the interrupt service routine desires to service the interrupting agent, but not completely perform the interrupt service, it can cause the external agent to negate the interrupt input but leave interrupt service pending through the use of the SW bits of the Cause register.

Catastrophic Exceptions

There are certain types of exceptions that indicate fundamental problems with the system. Although there is little the software can do to handle such events, they are worth discussing. Exceptions such as these are typically associated with faulty systems, such as in the initial debugging or development of the system.

Potential problems can arise because the processor does not automatically stack context information when an exception is detected. If the processor context has not been preserved when another exception is recognized, the value of the status, cause, and EPC registers are lost and thus the original task can not be resumed.

An example of this occurring is an exception handler performing a memory reference that results in a bus error (for example, when attempting to preserve context). The bus error forces execution to the exception vector location, overwriting the status, cause, and context registers. Proper operation cannot be resumed.

Handling Specific Exceptions

This section documents some specific issues and techniques for handling particular R36100 exceptions.

Address Error Exception

Cause

This exception occurs when an attempt is made to load, fetch, or store a word that is not aligned on a word boundary. Attempting to load or store a half-word that is not aligned on a half-word boundary will also cause this exception. The exception also occurs in User mode if a reference is made to a virtual address whose most significant bit is set (a kernel address). This exception is not maskable.

Handling

The R36100 branches to the General Exception vector for this exception. When the exception occurs, the CPU sets the ADEL or ADES code in the Cause register *ExcCode* field to indicate whether the address error occurred during an instruction fetch or a load operation (ADEL) or a store operation (ADES).

The EPC register points at the instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the EPC register points at the branch instruction that preceded the exception-causing instruction and sets the BD bit of the Cause register.

The R36100 saves the *KUp*, *IEp*, *KUc*, and *IEc* bits of the Status register in the *KUo*, *IEo*, *KUp*, and *IEp* bits, respectively and clears the *KUc* and *IEc* bits.

When this exception occurs, the *BadVAddr* register contains the virtual address that was not properly aligned or that improperly addressed kernel data while in User mode. The contents of the *VPN* field of the *Context* and *EntryHi* registers are undefined.

Servicing

A kernel should hand the executing process a segmentation violation signal. Such an error is usually fatal; although, an alignment error might be handled by simulating the instruction that caused the error.

Breakpoint Exception

Cause

This exception occurs when the R36100 executes the *BREAK* instruction and is not maskable.

Handling

The R36100 branches to the General Exception vector for the exception and sets the BP code in the *CAUSE* register *ExcCode* field.

The R36100 saves the *KUp*, *IEp*, *KUc*, and *IEc* bits of the Status register in the *KUo*, *KUp*, and *IEp* bits, respectively, and clears the *KUc* and *IEc* bits.

The *EPC* register points at the *BREAK* instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the *EPC* register points at the *BRANCH* instruction that preceded the *BREAK* instruction and sets the *BD* bit of the Cause register.

Service

The breakpoint exception is typically handled by a dedicated system routine. Unused bits of the BREAK instruction (bits 25..6) can be used pass additional information. To examine these bits, load the contents of the instruction pointed at by the *EPC* register.

Note: If the instruction resides in the branch delay slot, add four to the contents of the *EPC* register to find the instruction.

To resume execution, change the *EPC* register so that the R36100 does not execute the BREAK instruction again. To do this, add four to the *EPC* register before returning.

Note: If a BREAK instruction is in the branch delay slot, the BRANCH instruction must be interpreted in order to resume execution.

Bus Error Exception**Cause**

This exception occurs when the Bus Error input to the CPU is asserted by external logic during a read operation. For example, events like bus time-outs, backplane bus parity errors, and invalid physical memory addresses or access types can signal exception. This exception is not maskable.

This exception is used for synchronously occurring events such as cache miss refills. The general interrupt mechanism must be used to report a bus error that results from asynchronous events such as a buffered write transaction.

Handling

The R36100 branches to the General Exception vector for this exception. When exception occurs, the R36100 sets the *IBE* or *DBE* code in the CAUSE register ExcCode field to indicate whether the error occurred during an instruction fetch reference (*IBE*) or during a data load or store reference (*DBE*).

The *EPC* register points at the instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the *EPC* register points at the BRANCH instruction that preceded the exception-causing instruction and sets the BD bit of the cause register.

The R36100 saves the *KUp*, *IEp*, *KUc*, and *IEc* bits of the Status register in the *KUo*, *IEo*, *KUp*, and *IEp* bits, respectively, and clears the *KUc* and *IEc* bits.

Servicing

The physical address where the fault occurred can be computed from the information in the CPO registers:

- If the Cause register's *IBE* code is set (showing an instruction fetch reference), the virtual address resides in the *EPC* register.
- If the Cause register's *DBE* exception code is set (specifying a load or store reference), the instruction that caused the exception is at the virtual address contained in the *EPC* register (if the BD bit of the cause register is set, add four to the contents of the *EPC* register). Interpret the instruction to get the virtual address of the load or store reference and then use the TLBProbe (*tlbp*) instruction and read EntryLo to compute the physical page number.

A kernel should hand the executing process a bus error when this exception occurs. Such an error is usually fatal.

Co-processor Unusable Exception

Cause

This exception occurs due to an attempt to execute a co-processor instruction when the corresponding co-processor unit has not been marked usable (the appropriate CU bit in the status register has not been set). For CP0 instructions, this exception occurs when the unit has not been marked usable and the process is executing in User mode: CP0 is always usable from Kernel mode regardless of the setting of the CP0 bit in the status register. This exception is not maskable.

Handling

The R36100 branches to the General Exception vector for this exception. It sets the CPU code in the CAUSE register ExcCode field. Only one co-processor can fail at a time.

The contents of the cause register's CE (Co-processor Error) field show which of the four co-processors (3,2,1, or 0) the R36100 referenced when the exception occurred.

The EPC register points at the co-processor instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the EPC register points at the branch instruction that preceded the co-processor instruction and sets the BD bit of the Cause register.

The R36100 saves the KUp, IEp, KUc, and IEc bits of the status register in the KUo, IEo, KUp, and IEp bits, respectively, and clears the KUc and IEc bits.

Servicing

To identify the co-processor unit that was referenced, examine the contents of the Cause register's CE field. If the process is entitled to access, mark the co-processor usable and restore the corresponding user state to the co-processor.

If the process is entitled to access to the co-processor, but the co-processor is known not to exist or to have failed, the system could interpret the co-processor instruction. If the BD bit is set in the Cause register, the BRANCH instruction must be interpreted; then, the co-processor instruction could be emulated with the EPC register advanced past the co-processor instruction.

If the process is not entitled to access to the co-processor, the process executing at the time should be handed an illegal instruction/privileged instruction fault signal. Such an error is usually fatal.

Interrupt Exception

Cause

This exception occurs when one of eight interrupt conditions (software generates two, hardware generates six) occurs.

Each of the eight external interrupts can be individually masked by clearing the corresponding bit in the IntMask field of the status register. All eight of the interrupts can be masked at once by clearing the IEc bit in the status register.

Handling

The R36100 branches to the General Exception vector for this exception. The R36100 sets the INT code in the Cause register's ExcCode field.

The IP field in the Cause register show which of six external interrupts are pending, and the SW field in the cause register shows which two software interrupts are pending. More than one interrupt can be pending at a time.

The R36100 saves the KUp, IEp, KUC, and IEC bits of the status register in the KUo, IEo, KUp, and IEp bits, respectively, and clears the KUC and IEC bits.

Servicing

If software generates the interrupt, clear the interrupt condition by setting the corresponding Cause register bit (SW1:0) to zero.

If external hardware generated the interrupt, clear the interrupt condition by alleviating the conditions that assert the interrupt signal.

Overflow Exception**Cause**

This exception occurs when an **ADD ADDI, SUB, or SUBI** instruction results in two's complement overflow. This exception is not maskable.

Handling

The R36100 branches to the General Exception vector for this exception. The R36100 sets the OV code in the CAUSE register.

The EPC register points at the instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the EPC register points at the Branch instruction that preceded the exception-causing instruction and sets the BD bit of the CAUSE register.

The R36100 saves the KUp, IEp, KUC, and IEC bits of the status register in the KUo, IEo, KUp, and IEp bits, respectively, and clears the KUC and IEC bits.

Servicing

A kernel should hand the executing process a floating point exception or integer overflow error when this exception occurs. Such an error is usually fatal.

Reserved Instruction Exception**Cause**

This exception occurs when the R36100 executes an instruction whose major opcode (bits 31..26) is undefined or a Special instruction whose minor opcode (bits 5..0) is undefined.

This exception provides a way to interpret instructions that might be added to or removed from the MIPS processor architecture.

Handling

The R36100 branches to the General Exception vector for this exception. It sets the RI code of the Cause register's ExcCode field.

The EPC register points at the instruction that caused the exception, unless the instruction is in a branch delay slot: in that case, the EPC register points at the Branch instruction that preceded the reserved instruction and sets the BD bit of the CAUSE register.

The R36100 saves the KUp, IEp, KUC, and IEC bits of the status register in the KUo, IEo, KUp, and IEp bits, respectively, and clears the KUC and IEC bits.

Servicing

If instruction interpretation is not implemented, the kernel should hand the executing process an illegal instruction/reserved operand fault signal. Such an error is usually fatal.

An operating system can interpret the undefined instruction and pass control to a routine that implements the instruction in software. If the undefined instruction is in the branch delay slot, the routine that implements the instruction is responsible for simulating the branch instruction after the undefined instruction has been "executed". Simulation of the branch instruction includes determining if the conditions of the branch were met and transferring control to the branch target address (if required) or to the instruction following the delay slot if the branch is not taken. If the branch is not taken, the next instruction's address is $[EPC] + 8$. If the branch is taken, the branch target address is calculated as $[EPC] + 4 + (\text{Branch Offset} * 4)$.

Note that the target address is relative to the address of the instruction in the delay slot, not the address of the branch instruction. For details on how branch target addresses are calculated, refer to the description of branch instruction.

Reset Exception**Cause**

This exception occurs when the R36100 RESET signal is asserted and then de-asserted.

Handling

The R36100 provides a special exception vector for this exception. The Reset vector resides in the R36100's un-mapped and un-cached address space; Therefore the hardware need not initialize the Translation Lookaside Buffer (TLB) or the cache to handle this exception. The processor can fetch and execute instructions while the caches and virtual memory are in an undefined state.

The contents of all registers in the R36100 are undefined when this exception occurs except for the following:

- The SWc, KUc, and IEc bits of the Status register are cleared to zero.
- The BEV bit of the Status register is set to one.
- The TS bit of the Status register is frozen at one.
- The Config register is unlocked and initialized as described in Chapter 5.

Servicing

The reset exception is serviced by initializing all processor registers, co-processor registers, caches, and the memory system. Typically, diagnostics would then be executed and the operating system bootstrapped, including setting of the PortSize, Config, and BusCtrl registers. The reset exception vector is selected to appear in the uncached, un-mapped memory space of the machine so that instructions can be fetched and executed while the cache and virtual memory system are still in an undefined state.

System Call Exception

Cause

This exception occurs when the R36100 executes a SYSCALL instruction.

Handling

The R36100 branches to the General Exception vector for this exception and sets the SYS code in the CAUSE register's ExcCode field.

The EPC register points at the SYSCALL instruction that caused the exception, unless the SYSCALL instruction is in a branch delay slot: in that case, the EPC register points at the branch instruction that preceded the SYSCALL instruction and the BD bit of the CAUSE register is set.

The R36100 saves the KUp, IEp, KUC, and IEc bits of the status register in the KUo, IEo, KUp, and IEp bits, respectively, and clears the KUC and IEc bits.

Servicing

The operating system transfers control to the applicable system routine. To resume execution, alter the EPC register so that the SYSCALL instruction does not execute again. To do this, add four to the EPC register before returning.

Note: If a SYSCALL instruction is in a branch delay slot, the branch instruction must be interpreted in order to resume execution.

Introduction

The IDT R36100 RISController is an integration of system memory controllers and peripherals around an R30xx family core. Thus, the system interface can be described at many levels:

- Operation of the execution core, including caches and write buffers.
- Operation of the various memory controllers during external transactions.
- Operation during internal peripheral transactions.

This chapter will provide an overview of these interfaces, and a complete description of the rules for the internal core. Detail on each of the internal memory controllers and peripherals are found in the later chapters. This chapter also includes an overview of the write interface and detailed timing diagrams of the write interface.

Bus Interface Overview

The R36100 RISController bus interface uses a separate de-multiplexed address and data bus, along with control signals which select the targeted memory resource and perform the necessary data path steering. Figure 7.1 is a conceptual representation of the R36100 bus interface.

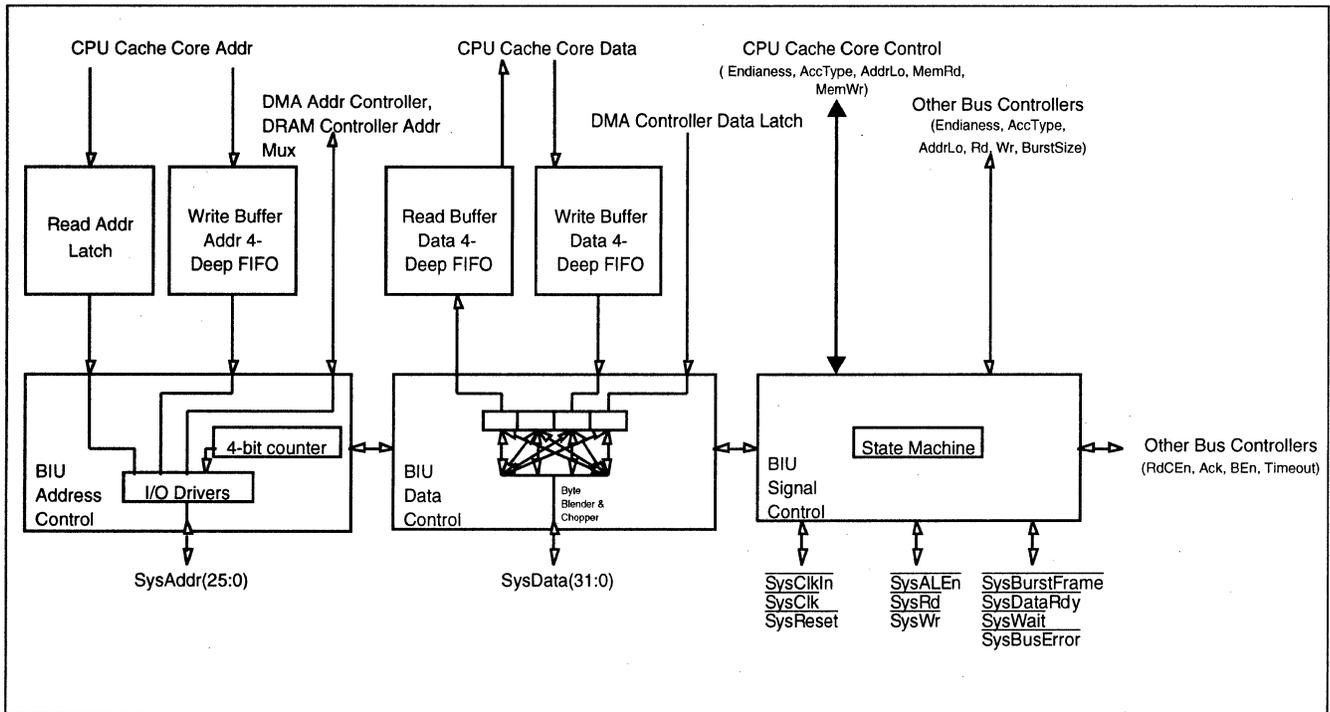


Figure 7.1 R36100 Bus Interface Unit Block Diagram

The address bus of the R36100 is a 26-bit address bus. Although the address bus is only 26 bits wide, the address space of the R36100 is actually 32 bits wide; the internal address decoder provides "Chip Selects" to target particular memory subsystems; thus, the width of the address bus only limits the size of any one memory sub-system, not the overall addressable memory.

The data bus is a 32-bit wide bus, with the capability to gather or parcel out data into smaller pieces when working with 8- or 16-bit data ports. Thus, the R36100 can mate directly with 8-, 16- and/or 32-bit memory subsystems; in fact, the widths of the various sub-systems are independently programmable through the various control registers. In addition, the R36100 can be used to implement either Big- or Little-endian memory systems, as selected at reset.

The control signals provided by the R36100 enable the processor to directly connect with a wide variety of external memory devices and peripherals. Wait-state generation and address decode is performed internally by the processor; once the proper external device for a transfer is determined by the R36100, the output control signals implement the protocol and timing selected for that memory sub-space.

During accesses to DRAM, the address bus will first carry the row address, then the column address, in a fashion suitable for direct connection with external page mode DRAM devices. Additional control signals provide $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, and transceiver control. Thus, the R36100 does not require external address multiplexors or complicated DRAM control state machines.

These techniques enable the R36100 to simply implement a wide variety of low cost systems, minimizing both system cost and development time, while maintaining high system throughput.

Pin Description

This section describes the signals used in the basic bus interface. Detailed information on the behavior of these signals, and of other signals used in other peripheral or memory control subsystems, are described in later chapters.

Note: Many R36100 signals have multiple functions, the exact behavior of a given pin is typically selected at device reset, and that signals indicated with an overbar (for example, $\overline{\text{overbar}}$) are active low.

System Bus Interface Signals

These signals are used by the bus interface to generate and provide global read and write signals.

SysAddr(25:0)

Output/(Input during external DMA)

System Address Bus: The R36100 uses a dedicated 26-bit physical address bus always driven by the R36100, except during the Address Strobe portion of external DMA cycles.

The address first becomes valid on the same (first) clock cycle that the address latch enable indicator, SysALEn, asserts. Coincidentally either SysRd or SysWr will also assert on the first valid cycle. The address is valid until either SysRd or SysWr de-asserts at the end of a transaction.

During Memory or I/O transactions, SysAddr(25:0) contain the internally latched 26-bit physical address. The SysAddr(3:0) bits represent the doubleword, word, halfword, and byte addresses that count with each datum during quad-word-burst and mini-burst reads and writes.

During DRAM transactions, SysAddr(13:2) are driven with the multiplexed DRAM row and column address. The SysAddr(4:2) bits during the column address period represent the doubleword, word, or halfword addresses that count with each datum during quad-word-burst and mini-burst reads and writes.

During idle cycles between valid transactions, the behavior of Addr(25:0) is undefined.

SysData(31:0)**Input/Output**

System Data Bus: The R36100 uses a dedicated 32-bit data bus. For reads, data is sampled on the rising edge of the $\overline{\text{SysClk}}$ reference clock. Although the data bus is 32-bits wide, the R36100 directly supports the use of narrower memory subsystems. In these cases, the R36100 bus unit interface will gather smaller data into the requested transfer size on reads, and break write data up into a series of smaller pieces on writes, depending on whether the port is 32, 16 or 8-bits wide. The bus interface will shift and adjust the LSB SysAddr bits accordingly for big or little endian data. Collectively, when these types of transactions are a result of word size or smaller accesses these are referred to as "mini-bursts" and quad-word accesses are referred to as "bursts". However, on the R36100, in contrast to the R3051 family, both mini-bursts and bursts assert the system burst signal, SysBurstFrame.

Note: During internal peripheral register reads, SysData(15:0) is driven with the register contents by the R36100. Also note that the SysData Bus may tri-state when the Bus Interface Unit is not in use.

Clock and Reset Signals**SysClkIn****Input**

System Clock Input: This is a double frequency input clock used to generate the timing of the processor.

 $\overline{\text{SysClk}}$ **Output**

System Output Clock: This output clock provides the master timing reference for all bus interface signals. All input signals are sampled on the rising edge of SysClk, and all outputs except for the external strobes $\overline{\text{DramCAS}}(3:0)$ and $\overline{\text{SysWrEn}}(3:0)$ are generated from the rising edge of SysClk. Thus, external logic can use the rising edge of the $\overline{\text{SysClk}}$ output reference to generate control signals back to the R36100, and to sample R36100 outputs.

The number of loads on $\overline{\text{SysClk}}$ should be kept below or equal to a maximum of 5 CMOS type loads due to internal pin skew feedback monitoring on the R36100. Systems requiring additional loads should either buffer or invert SysClk.

There is no guaranteed AC timing delay relationship between the SysClkIn input clock and the $\overline{\text{SysClk}}$ output clock. However, the phase relationship can be guaranteed via the method described in Chapter 19, "Debug Mode Features."

 $\overline{\text{Reset}}$ **Input**

Reset: This active low input signal initializes the processor and is required after power up before correct operation can begin. Optional features of the processor including the Endianess and the Port Width of the Boot ROM are established during the last cycle of reset using the reset configuration mode inputs (also known as the reset initialization vector) which are multiplexed with the interrupt pins. See Chapter 18 for more information on the reset initialization vector.

Bus Interface Control Signals **$\overline{\text{SysALEn}}$** **Output (Input during external DMA)**

System Address Latch Enable: As an output signal, this active low signal indicates when a new address is first valid. $\overline{\text{SysALEn}}$ de-asserts high one clock after being asserted by the R36100.

As an input signal, it is used by an external DMA agent to indicate that it has provided a valid address on the processor data bus. The R36100 will then use this address to select which memory subsystem or peripheral is the target of the DMA, and perform the necessary access. As an input, its timing is similar to the output case; however, its assertion by the external DMA agent may be 1 or more clocks long. The address is sampled by the CPU on the first rising $\overline{\text{SysClk}}$ edge where $\overline{\text{SysALEn}}$ is asserted. Any additional $\overline{\text{SysALEn}}$ asserted clocks are ignored as far as the address is concerned, however, its assertion delays the second phase of the DMA transaction where the external DMA agent tri-states most of the bus control signals and lets the R36100 memory controllers drive the bus control signals.

 $\overline{\text{SysRd}}$ **Output (Input during external DMA)**

System Read: $\overline{\text{SysRd}}$ is always driven by the R36100, except during external DMA cycles.

As an output, $\overline{\text{SysRd}}$ is an active low read control signal. During external read transfer cycles, this signal will be asserted. This signal can be used for external control and diagnostics as needed.

Note: During internal peripheral register reads, the SysData(15:0) is driven with the register contents by the R36100.

As an input, this signal can be driven by an external DMA engine. However, the R36100 does not use this signal during DMA. If $\overline{\text{SysWr}}$ is high when $\overline{\text{SysALEn}}$ is asserted, then a DMA read is implied.

 $\overline{\text{SysWr}}$ **Output (Input during external DMA)**

System Write: $\overline{\text{SysWr}}$ is always driven by the R36100, except during external DMA cycles.

As an output, $\overline{\text{SysWr}}$ is an active low control signal to indicate that the current transaction is a write. This signal can be used for external control and diagnostics as needed.

As an input, this signal is used internally during DMA as a read/write input signal. If $\overline{\text{SysWr}}$ is high when $\overline{\text{SysALEn}}$ is de-asserted, then a DMA read is implied.

 $\overline{\text{SysBurstFrame}}$ **Output (Input during external DMA)**

System Burst: $\overline{\text{SysBurstFrame}}$ is always driven by the R36100, except during external DMA transactions.

This active low control signal specifies that at least one more datum will be written or read. This signal is valid for both reads and writes. $\overline{\text{SysBurstFrame}}$ always asserts on the first clock of a transaction. Thereafter it de-asserts high on the first clock of the last datum. Thus, if there is only a single datum, $\overline{\text{SysBurstFrame}}$ will assert for one clock only. If there are multiple datum, $\overline{\text{SysBurstFrame}}$ remains asserted until the last datum begins.

However, an external DMA has the ability to continue a burst for an indefinite length. Under this condition, the R36100 is unable to determine when the last datum will occur, and it is assumed that the external DMA agent is able to determine how many datum it needs to read or write.

So, on external DMA transactions, $\overline{\text{SysBurstFrame}}$ is an input signal for the first clock cycle of an R36100 DRAM, memory, or I/O controller access. During the bus transaction phase of external DMA, $\overline{\text{SysBurstFrame}}$ reverts back to an output, but is undefined after the $\overline{\text{SysALE}}$ phase.

SysDataRdy**Output**

System Data Ready: This active low output signal indicates that the CPU is ready to receive data on a read or that it is driving data on a write.

If wait-states are inserted by the use of SysWait (as opposed to the internally wait-state generator), then SysDataRdy is kept asserted for the same number of clock cycles. Thus external diagnostic tools such as logic analyzers may want to gate SysDataRdy with SysWait.

During external DMA transactions, SysDataRdy is still driven by the CPU the entire time. Thus during the second part of the DMA transaction where the R36100 memory controller is being used to access data, the CPU drives SysDataRdy whenever it has or expects valid data.

SysWait**Input**

System Wait: Active low is used to extend the length of the memory cycle by stalling the Memory, or I/O Controller for as many cycles as SysWait is sampled low. SysWait can be asserted anytime during single word or smaller read and write transactions. However, the effect of SysWait occurs 1 clock later due to internal pipelining of the signal. During burst reads, SysWait is valid up until the internal Acknowledge is generated to release the execution core for refill cycles. Thus systems that are quad-word bursting and using SysWait must either not use the internal Acknowledge, or guarantee that SysWait is only asserted during the early part of a cycle with the Memory or I/O Controller Wait option programmed.

The specific behavior of SysWait is dependent on the type of memory accessed. In general, its effect is to delay the de-assertion of the pertinent data strobe, such as MemRdEnOdd.

SysBusError**Input**

System Bus Error: This is an active low input signal which terminates a bus transaction on the next clock. If an internal Acknowledge has not already been generated on a read access, a bus error exception will be generated. If on a read access, SysBusError is asserted after the internal Acknowledge, then the bus transaction will be terminated, but no error will be reported. SysBusError assertion during a write access always end the bus transaction on the next clock; however, a bus error exception is never generated for a write access.

CPU Core Transaction Types

The R36100 RISController execution core is capable of requesting the following types of transactions:

Read Operation

The processor executes an instruction fetch or a data load operation as the result of either a cache miss or an uncacheable reference.

Quad word reads occur when the processor requests a contiguous block of four words from memory. Quad word reads occur in response to instruction cache misses, and will occur in response to a data cache miss if the DBlockRefill option in the CPO Cache Configuration register is enabled. The R36100 incorporates an on-chip 4-word deep read buffer which may be used to "queue up" the read response before passing it through to the high-bandwidth cache and execution core. Read buffering is appropriate in systems which require wait states between adjacent datums of a block read or in interfacing to memory systems narrower than 32-bits wide.

On the other hand, systems that use high-bandwidth memory techniques—such as page mode, static column, nibble mode, or memory interleaving—can effectively bypass the read buffer by providing words of the block at the processor clock rate. Note that the choice of burst vs. read buffering is independent of the initial latency of the memory; that is, burst mode can be used even if multiple wait states are required to access the first datum of the block.

Single data reads (single word, tri-byte, halfword, or byte) are used for uncacheable references (such as for I/O or boot code) and will be used in response to a data cache miss if the DBlockRefill option in the CP0 Cache Configuration register is disabled. A single data read returns one unit of data per read transaction.

Write Operations

The R36100 utilizes an on-chip write buffer to isolate the execution core from the speed of external memory during write operations. There is a single primary type of write:

Single data writes (word, tri-byte, halfword, or byte writes corresponding to 32-bit, 16-bit, and 8-bit interfaces, respectively) are used in response to a store operation, either cached or uncached (the R36100 uses a write-through cache).

Although the CPU execution core is capable of producing only single data writes, the DMA Controller can produce 4 word burst writes.

Quad word writes occur when the DMA Controller is instructed to transfer 4 words at a time. The DMA Controller will first read 4 words into the read buffer, then latch out the data for 4 consecutive writes.

Although the data bus is 32-bits wide, the R36100 directly supports the use of narrower memory subsystems. In these cases, the R36100 will gather smaller data into the requested transfer size on reads, and break write data up into a series of smaller pieces on writes. Collectively, these types of transactions are referred to as "mini-bursts".

Multiple Operations

It is possible for the R36100 execution core to have multiple activities pending. Specifically, there may be data in the write buffer, a read request—such as due to a cache miss—or two read requests (both the I-cache and D-cache misses in a single clock cycle), even as the bus interface is servicing some external DMA activity.

In establishing the order in which the requests are processed, the R36100 is sensitive to possible conflicts and data coherency issues as well as to performance issues. For example, if the on-chip write buffer contains data which has not yet been written to memory, and the processor issues a read request to the target address of one of the write buffer entries, then the processor strategy must insure that the read request is satisfied by the new, current value of the data.

There are two levels of priority: that performed by the CPU engine, and that performed by the bus interface unit. The internal execution engine can be viewed as making requests to the bus interface unit. In the case of multiple requests in the same clock cycle, the CPU core will:

- Perform the data request first. That is, if both the data cache and instruction cache miss in the same clock cycle, the processor core will request a read to satisfy the data cache first. Similarly, a write buffer full stall will be processed before an instruction cache miss.
- Perform a read due to an instruction cache miss.

This prioritization is important in maintaining the precise exception model of the MIPS architecture. Since data references are the result of instructions which entered the pipeline earlier, they must be processed (and any exceptions serviced) before subsequent instructions (and their exceptions) are serviced.

Once the processor core internally decides which type of request to make to the bus interface unit, it then presents that request to the bus interface unit.

In the R36100 Bus Interface Unit, multiple operations are serviced in the following order:

1. DRAM refreshes may delay the start of a read or write DRAM data access.
2. Ongoing transactions are completed without interruption.
3. DMA requests are serviced according to the DMA priorities established in the R36100. (DMA requests using the cache, stall the CPU and flush the write buffer before beginning).
4. Instruction cache misses are processed.
5. Pending writes are processed.
6. Data cache misses or uncacheable reads/uncacheable instruction fetches are processed.

This service order has been designed to achieve maximum performance, minimize complexity, and solve the data coherency problem possible in write buffer systems.

This order assumes that the write buffer does not contain instructions which the processor may wish to execute. The processor does not write directly into the instruction cache: store instructions generate data writes which may change only the data cache and main memory. The only way in which an instruction reference may reside in the write buffer is in the case of self modifying code, generated with the caches swapped. However, in order to unswap the caches, an uncacheable instruction which modifies CPO must be executed; the fetch of this instruction would cause the write buffer to be flushed to memory. Thus, this ordering enforces strong ordering of operations in hardware, even for self modifying code. Of course, software could perform an uncacheable reference to flush the write buffer at any time, thus achieving explicit memory synchronization with software.

Execution Engine Fundamentals

This section describes the fundamentals of the processor interface and its interaction with the execution core. These fundamentals will help to explain the relationship between design trade-offs in the system interface and the performance achieved in R36100 systems.

Execution Core Cycles

The R36100 execution core utilizes many of the same operation fundamentals as does the R3000A processor. Thus, much of the terminology used to describe the activity of the R36100 is derived from the terminology used to describe the R3000A. In many instances, the activity of the execution core is independent of that of the bus interface unit.

Cycles

A cycle is the basic timing reference of the R36100 execution core. Cycles in which forward progress is made (the processor pipeline advances) are called Run cycles. Cycles in which no forward progress occurs are called Stall cycles. Stall cycles are used for resolving exigencies such as cache misses, write stalls, and other types of events. All cycles can be classified as either run or stall cycles.

Run Cycles

Run cycles are characterized by the transfer of an instruction into the processor execution core, and the optional transfer of data into or out of the execution core. Thus, each run cycle can be thought of as having an instruction and data, or ID, pair.

There are two types of run cycles: cache-run cycles and refill-run cycles. Cache-run cycles, typically referred to as just-run cycles, occur while the execution core is executing out of its on-chip cache; these are the principal execution mechanism.

Refill-run cycles, referred to as streaming cycles, occur when the execution core is executing instructions as they are brought into the on-chip cache. For the R36100, streaming cycles are defined as cycles in which data is brought out of the on-chip read buffer into the execution core, rather than defining them as cycles in which data is brought from the memory interface to the read buffer.

Stall Cycles

There are three types of stall cycles:

Wait Stall Cycles. These are commonly referred to simply as stall cycles. During wait stall cycles, the execution core maintains a state consistent with resolving a stall causing event. No cache activity will occur during wait stalls.

Refill Stall Cycles. These occur only during memory reads, and are used to transfer data from the on-chip read buffer into the caches.

Fixup Stall Cycles. Fixup cycles occur during the final cycle of a stall; that is, one cycle before entering a run cycle or entering another stall. During the final fixup cycle (the one which occurs before finally re-entering run operation), the Instruction/Data (ID) pair which should have been processed during the last run cycle is handled by the processor. The fixup cycle is used to restart the processor and co-processor pipelines, and in general to fixup conditions which caused the stall.

There are five basic stalls that are caused by the following conditions:

Read Busy Stalls: If the processor core requires read data, either to process a cache miss or an uncacheable reference, then it will be stalled until the read data is brought back to the execution core.

Write Busy Stalls: If the processor attempts to perform a store operation while the on-chip write buffer is already full, then the processor will stall until a write transaction is begun on the interface to free up room in the write buffer for the new address and data.

Multiply/Divide Busy Stalls: If software attempts to read the result registers of the integer multiply/divide unit (the HI and LO registers) while a multiply or divide operation is underway, the processor execution core will stall until the results are available.

Micro-TLB¹ Fill Stalls: These stalls can occur when an instruction translation misses in the instruction TLB cache (the micro-TLB, which is a two-entry cache of the main TLB used to translate instruction references). When such an event occurs, the execution core will stall for one cycle, in order to refill the micro-TLB from the main TLB. Since this is a single-cycle stall, it is of necessity a fixup cycle.

¹ Micro-TLB stalls will not occur in the R36100, which does not include an on-chip TLB.

Multiple Stalls: Multiple stalls are possible whenever more than one stall initiating event occurs within a single run cycle. An example of this condition is when a single cycle results in both an instruction cache miss and a data cache miss. The most important characteristic of any multiple stall cycle is the validity of the Instruction/Data (ID) pair processed in the final fixup cycle. The R36100 execution core keeps track of nested stalls to insure that orderly operation is resumed once all of the stall causing events are processed.

For the general case of multiple stalls, the service order is:

1. Micro-TLB Miss and Partial Word Store
2. Data Cache Miss or Write Busy Stall
3. Instruction Cache Miss
4. Multiply/Divide Unit Busy

Internal Acknowledgment

To speed performance, the R36100 CPU core allows the CPU to exit wait stalls and begin refill and/or streaming, even while the bus interface continues to provide additional data to the CPU.

To do this, the R36100 incorporates an on-chip 4-entry read buffer. In response to a quad word read, data begins to be returned to the CPU. As each datum is returned, it is entered into the read buffer. At some point, the internal core is "Acknowledged" (the "AckN" internal signal) to indicate that the read buffer contents may begin being transferred to the internal caches and execution core.

Transfer from the read buffer to the core/caches occurs at the pipeline rate. Thus, the ideal time to provide such an acknowledgment is 3 cycles before the last datum is returned to the R36100. In this case, the last datum will be entered into the read buffer, and in the very next clock cycle be placed into the cache/core. Note that in the case of single word reads, acknowledge is provided with the last byte of the requested transfer; in the very next clock cycle, the datum is transferred into the core.

To facilitate this operation, the R36100 requires that the various memory controllers be programmed for the optimal placement of "Ack", the internal control signal which is used to begin refill/streaming. As a rule, Ack should be placed 3 cycles before the last response datum in a quad word read.

Read Interface Timing Overview

The read interface is designed to allow a variety of memory strategies. An overview of how data is transmitted from memory and I/O devices to the processor is discussed below.

Initiation of a Read Request

A read transaction occurs when the processor internally performs a run cycle which is not satisfied by the internal caches. Immediately after the run cycle, the processor enters a stall cycle and asserts the internal control signal $\overline{\text{MemRd}}$. This signals to the internal bus interface unit arbiter that a read transaction is pending.

Assuming that the read transaction can be immediately processed (that is, there are no ongoing bus operations and no higher priority operations pending), the processor will initiate a bus read transaction on the rising edge of SysClk which occurs during phase two of the processor stall cycle. Higher priority operations would have the effect of delaying the start of the read by inserting additional processor stall cycles.

Figure 7.2 illustrates the initiation of a read transaction, based on the internal assertion of the $\overline{\text{MemRd}}$ control signal. This figure is useful in determining the overall latency of cache misses on processor operation.

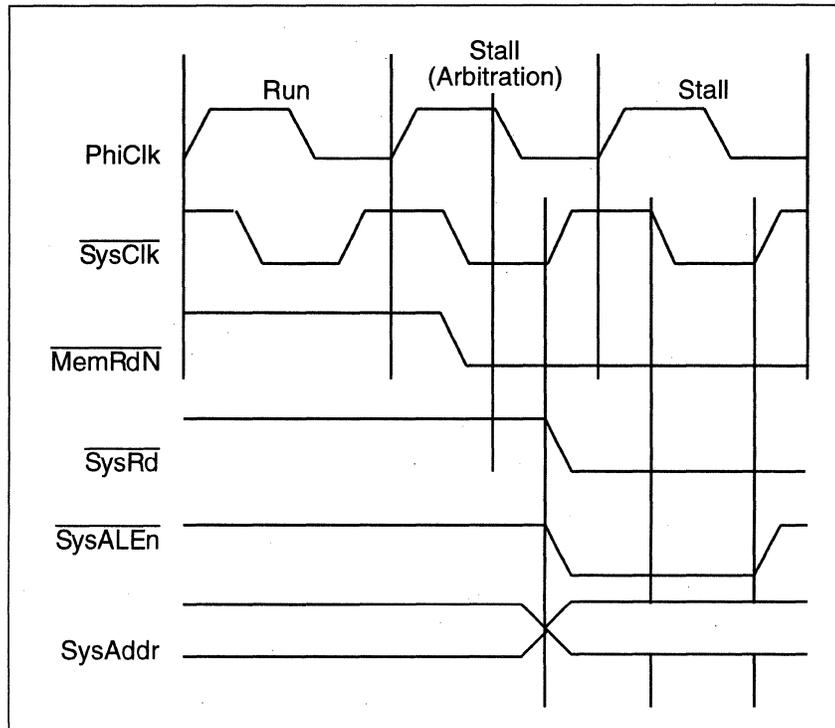


Figure 7.2 CPU Latency to Start of Read

Memory Addressing

A read transaction begins when the processor asserts its $\overline{\text{SysRd}}$ control output, and also drives the address and other control information onto the SysAddr and memory interface buses. Figure 7.3 illustrates the start of a processor read transaction.

The addressing occurs throughout the read transaction. At the rising edge of $\overline{\text{SysClk}}$, the processor will drive the read target address onto the SysAddr bus. At this time, $\overline{\text{SysALEn}}$ will also be asserted, to allow an external ASIC or peripheral to capture the address. During the initial part of the read phase, all-memory control read enables will be held high indicating that memory drivers should not be enabled onto the SysData bus.

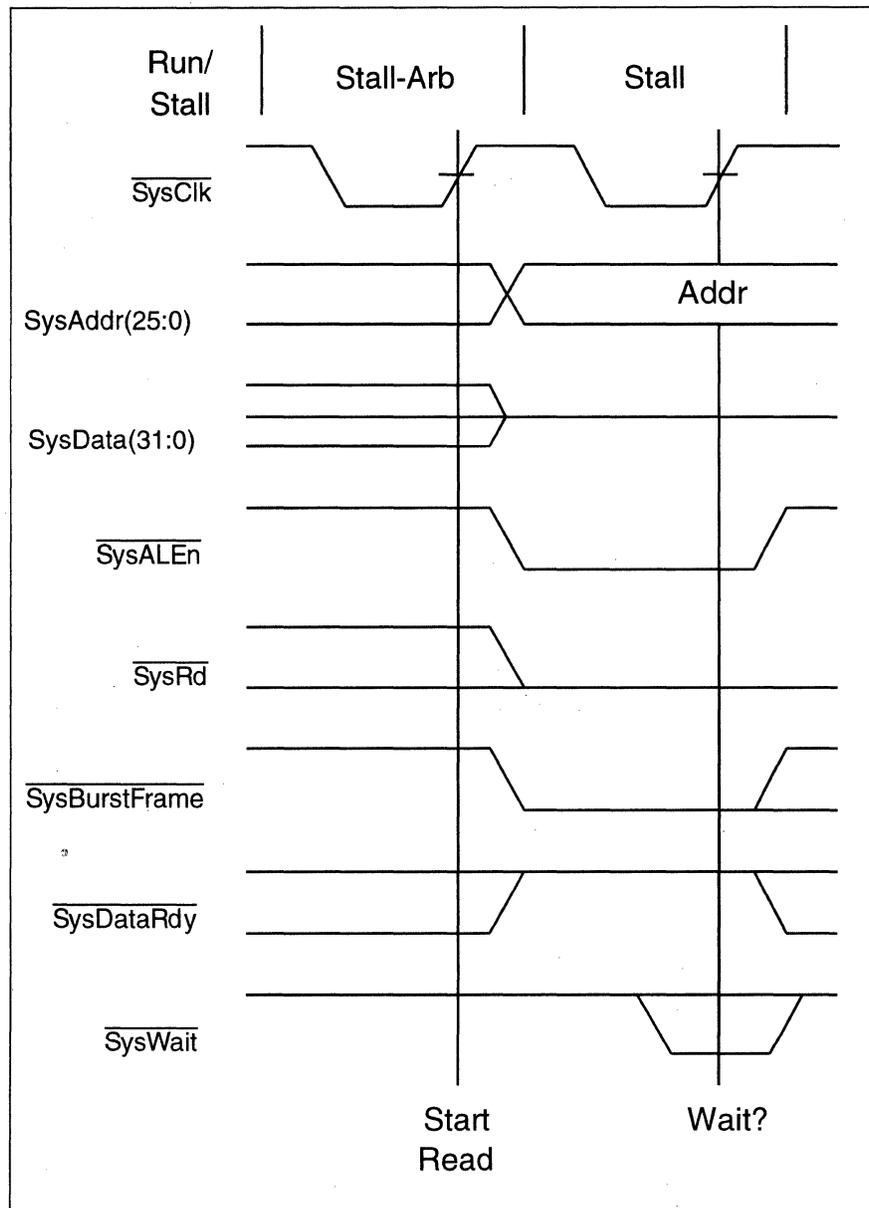


Figure 7.3 Start of Bus Read Operation

Concurrent with driving addresses on the SysAddr bus, the processor will redundantly indicate the beginning of the read transaction with SysBurstFrame asserting. A multi-datum transaction and bursts will be indicated by SysBurstFrame remaining asserted as the current datum is sampled. The functioning of the SysAddr(3:0) counter during mini-burst and burst reads is also described later.

Initiation of the Data Phase

Once the SysAddr bus has presented the address for the transfer, the various memory controller read enables assert and data is ready to be sampled.

Bringing Data into the Processor

Regardless of whether the transfer is a burst read or a single datum transfer, the basic mechanism for transferring data presented on the A/D bus into the processor is the same.

Although there are two internal control signals involved in terminating read operations, only the internal RdCEnN signal is used to cause data to be captured from the bus.

The memory system asserts internal RdCEnN to indicate to the processor that it has (or will have) data on the data bus to be sampled. The earliest that internal RdCEnN can be detected by the processor is the rising edge of $\overline{\text{SysClk}}$ after it has asserted $\overline{\text{SysALEn}}$ (start of phase 1 of the second clock cycle of the read).

If internal RdCEnN is detected as asserted by the internal wait-state generator, the processor will capture (with proper setup and hold time) the contents of the SysData bus on the immediately subsequent rising edge of $\overline{\text{SysClk}}$. This captures the data in the internal read buffer for later processing by the execution core/cache subsystem.

The R36100 integrates on-chip a 4-word read buffer, capable of acting as a speed-matching FIFO between the system interface and the execution core. This bus interface then performs byte or half-word gathering, and assembles them into 32-bit words for the read buffer. Thus, the bus interface supports 8-, 16-, and 32-bit memory subsystems, even for quad word reads, with no real system impact.

Figure 7.4 illustrates the sampling of data by the R36100.

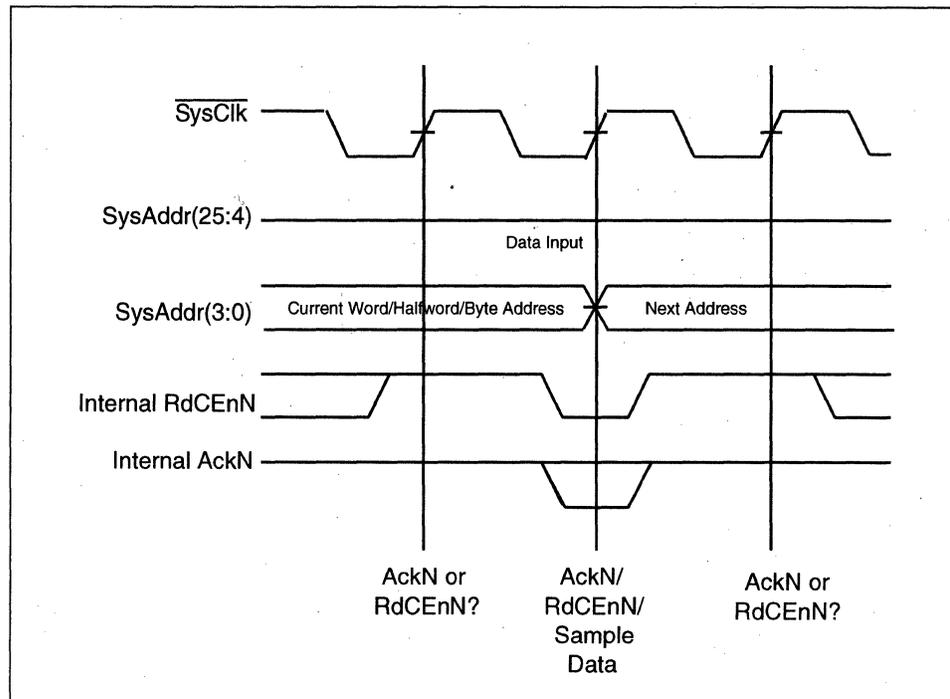


Figure 7.4 Data Sampling.

Terminating the Read

Following are three methods for the external memory system to terminate an ongoing read operation:

- It can supply an internal AckN (acknowledge) to the processor, to indicate that it has sufficiently processed the read request and has or will supply the requested data in a timely fashion. Note that internal AckN may be signalled to the processor “early”, to enable it to begin processing the read data even while additional data is brought from the SysData bus. This is applicable only in quad-word read operations.
- It can supply a $\overline{\text{SysBusError}}$ to the processor, to indicate that the requested data transfer has “failed” on the bus, and force the processor to take a bus error exception. Although the system interface behavior of the processor when $\overline{\text{SysBusError}}$ is presented is similar to the behavior when internal AckN is presented, no data will actually be written into the on-chip cache. Rather, the cache line will either remain unchanged, or will be invalidated by the processor, depending on how much of the read has already been processed. Thus it is recommended that if $\overline{\text{SysBusError}}$ is to be used for quad word burst transactions, that it be asserted on the first clock of the transaction.
- The external memory system can supply the requested data, using internal RdCEnN to enable the processor to capture data from the bus. The processor will “count” the number of times internal RdCEnN is sampled as asserted; once the processor counts that the memory system has returned the desired amount of data (one byte to four words), it will implicitly “acknowledge” the read after it samples the last required internal RdCEnN. This technique may be important in memory systems where the latency can vary--e.g. dual ported memory.

Throughout this chapter, method one will be illustrated. The other cases can be extrapolated easily from these diagrams (for example, the system designer can assume that internal AckN is asserted simultaneous with the last internal RdCEnN of a single word read transfer and 3 clocks before the last internal RdCEnN of a quad word burst read transfer).

There are actually two phases of terminating the read: there is the phase where the memory system indicates to the processor that it has sufficiently processed the read request, and the internal read buffer can be released to begin refilling the internal caches; and there is the phase in which the read control signals are negated by the processor bus interface unit.

The difference between these phases is due to block refill: it is possible for the memory system to “release” the execution core even though additional words of the block are still required; in that case, the processor will continue to assert the external read control signals until all four words are brought into the read buffer, while simultaneously refilling/executing based on the data already brought on board.

To determine the end of the read transaction, one of the following methods may be used:

- Systems that only use 32-bit memory sub-region ports as with the R3051 family only have single datum reads or burst reads and can either count the number of wait-cycles or use the de-asserting edge of $\overline{\text{SysRd}}$ to end the transaction.
- Systems that use 16 or 8-bit ports must in general support mini-burst (multi-datum) reads. Memory controllers for such systems can use the de-asserting edge of $\overline{\text{SysRd}}$ to reset the controller. The memory controller can also look for $\overline{\text{SysBurstFrame}}$ to de-assert. When $\overline{\text{SysBurstFrame}}$ de-asserts, the controller knows that it is handling the final datum of the transaction.

Figure 7.5 shows the timing of the control signals when the read cycle is being terminated.

Latency Between Processor Operations

In general, the processor may begin a new bus activity as soon as the phase immediately after the termination of the read cycle. Although this operation may logically be either a read, write, or bus grant, there are no cases where a read operation can be signalled by the internal execution core at this time.

Since a new operation may begin one-half clock cycle after the data is sampled from the bus, it is important that the external memory system cease to drive the bus prior to this clock edge. To simplify design, the processor provides various read enable outputs for each memory controller, which can be used to control either the Output Enable of the memory device (presuming its tri-state time is fast enough), or to control the Output Enable of a buffer or transceiver between the memory device data bus and the processor SysData bus.

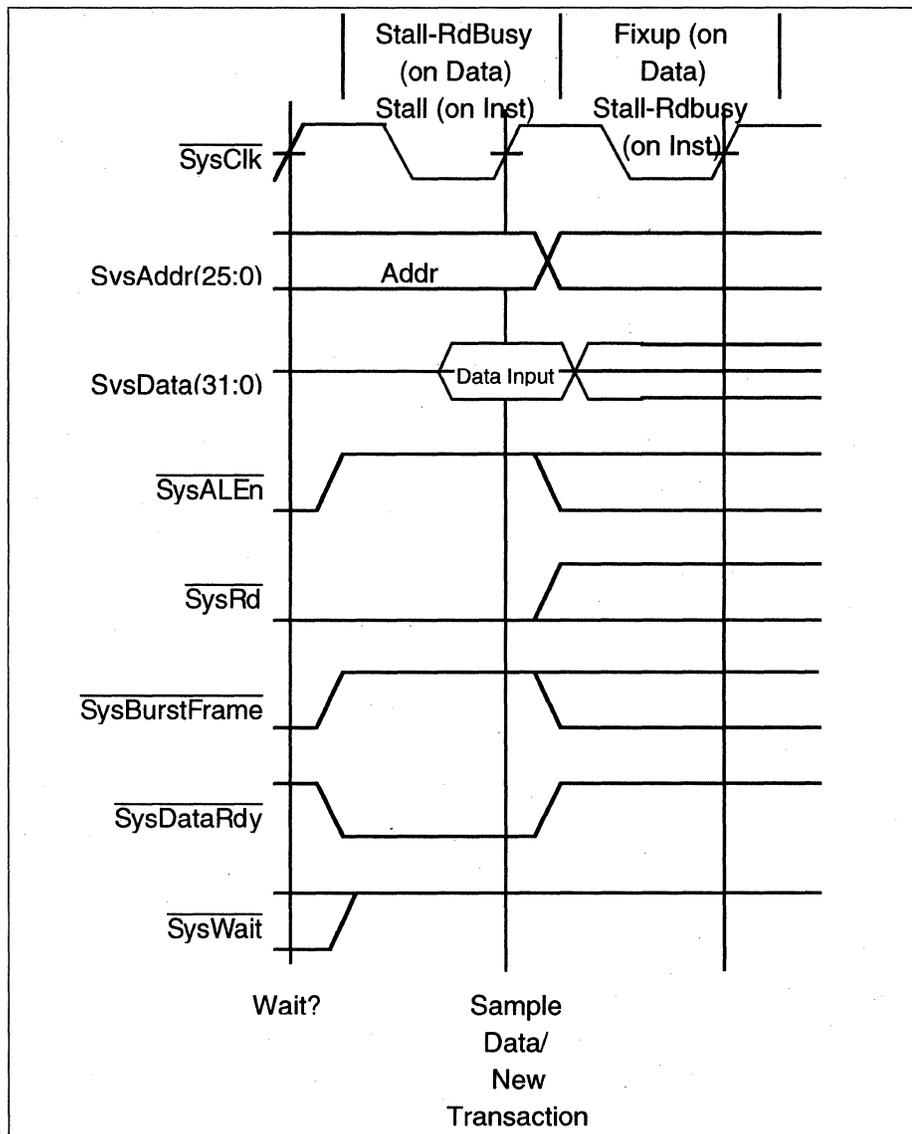


Figure 7.5 Read Cycle Termination.

The R36100 also adds a feature to the R3051 family to enable the system designer to lengthen the amount of time available for bus turn-around. The Bus Turn Around control field of the various memory controller Control Registers enables the system designer to extend the minimum guaranteed amount of time available for bus turn-around. This enables the system designer to eliminate some transceiver devices and/or use slower system components, without worrying about bus conflicts.

Processor Internal Activity

In general, the processor will execute stall cycles until an internal AckN is detected. It will then begin the process of refilling the internal caches from the read buffer.

The system designer should consider the difference between the time when the memory interface has completed the read, and when the processor core has completed the read. The bus interface may have successfully returned all of the required data, but the processor core may still require additional clock cycles to bring the data out of the read buffer and into the caches. Figure 7.6 illustrates the relationship between AckN and the internal activity for a block read.

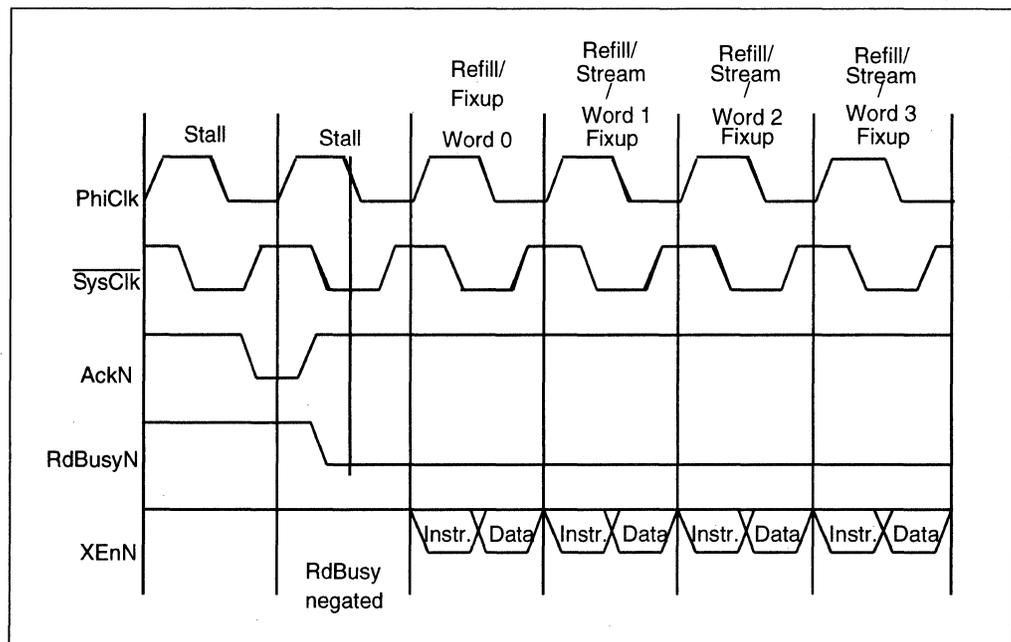


Figure 7.6 Internal Processor States on 4-word Burst Read.

This figure illustrates that the processor may perform either a stream, fixup, or refill cycle in cycles in which data is brought from the read buffer. The difference between these cycles is defined as follows:

- **Refill.** A refill cycle is a clock cycle in which data is brought out of the read buffer and placed into the internal processor cache. The processor does not execute on this data.
- **Fixup.** A fixup cycle is a cycle in which the processor transitions into executing the incoming data. It can be thought of as a “retry” of the cache cycle which resulted in a miss.
- **Stream.** A stream cycle is a cycle in which the processor simultaneously refills the internal instruction cache and executes the instruction brought out of the read buffer.

When reading the block from the read buffer, the processor will use the following rules:

- For uncacheable references, the processor will bring the single word out of the read buffer using a fixup cycle.
- For data cache refill, it will execute either one or four refill cycles, followed by a fixup cycle.
- For instruction cache refill, it will execute refill cycles starting at word zero until it encounters the miss address, and then transition to a fixup cycle. It will then execute stream cycles until either the entire block is processed, or an event stops execution. If something causes execution to stop, the processor will process the remainder of the block using simple refill cycles. For example, Figure 7.7 illustrates the refill/fixup/stream sequence appropriate for a miss which occurs on the second word of the block (word address 1).

Although this operation is transparent to the external memory system, it is important to understand this operation in order to gauge the impact of design trade-offs on performance.

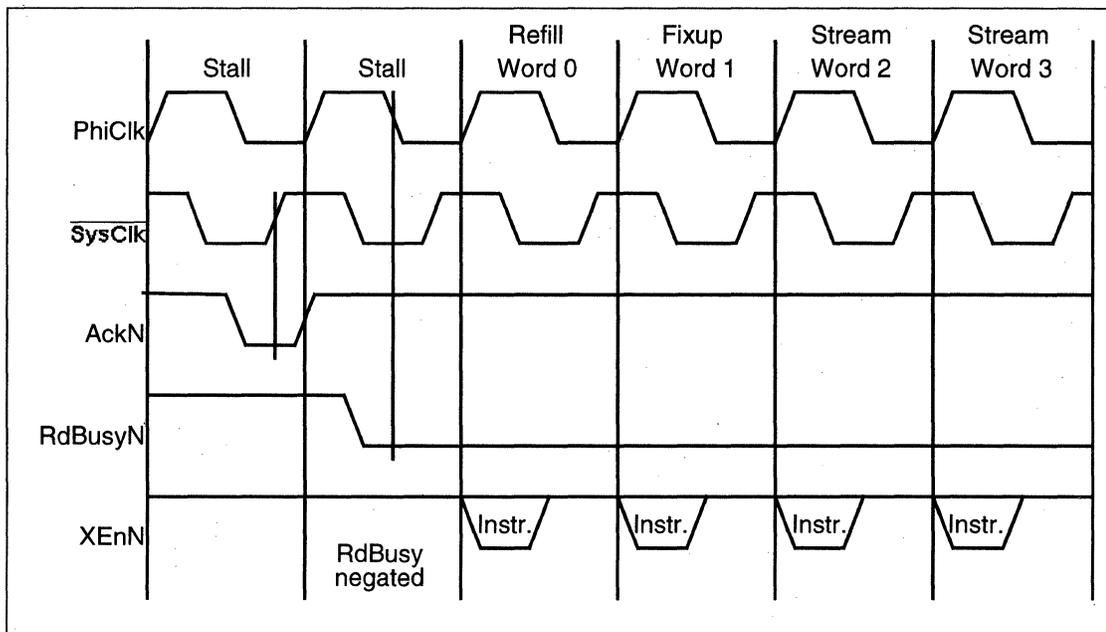


Figure 7.7 Instruction Streaming Internal Operation Example.

The Write Interface

The write protocol of the R36100 has been designed to complement the read interface of the processor. Many of the same signals are used for both reads and writes, simplifying the design of the memory system control logic.

Importance of Writes in R36100 Systems

The design goal of the write interface is to insure that a relatively slow write cycle does not degrade the performance of the processor. To this end, a four deep write buffer has been incorporated on-chip. The role of the write buffer is to decouple the speed of the memory interfaces from the speed of the execution engine.

The write buffer captures store information (data, address, and transaction size) from the processor at its clock rate, and later presents it to the memory interface at the rate it can perform the writes. Four such buffer entries are incorporated, thus allowing the processor to continue execution even when performing a quick succession of writes. Only when the write buffer is already filled will the processor stall; simulations have shown that significantly less than 1% of processor clock cycles are lost to write buffer full stalls.

Although it may be counter-intuitive, a significant percentage of the bus traffic will in fact be processor writes to memory. This can be demonstrated if one assumes the following:

Instruction Mix:

ALU Operations	55%
Branch Operations	15%
Load Operations	20%
Store Operations	10%

Cache Performance:

Instruction Hit Rate	95%
Data Hit Rate	90%

For these assumptions, in 100 instructions, the bus would see:

- 5 Reads to process instruction cache misses on instruction fetches
- $10\% \times 20 = 2$ reads to process data cache misses on loads
- 10 store operations to the write through cache
- Total: 7 reads and 10 writes

In this example, about 60% of the bus transactions are write operations, even though only 10 instructions were store operations versus 100 instruction fetches and 20 data fetches.

Types of Write Transactions

The R36100 has two basic types of write transactions, depending on the port size selected in the CP0 Port Size Configuration register for each memory sub-region. When writes are generated from the CPU core, the 32-bit ports use only the single-word write type. DMA channels are also able to generate burst writes. The 16-bit ports can use the single half-word write or the mini-burst (double halfword) write type. And 8-bit ports are able to use the single byte write or the mini-burst (double, tri, or quad byte) or DMA burst write types.

32-Bit Write Transactions

Unlike instruction fetches and data loads, which are usually satisfied by the on-chip caches and thus are not seen at the bus interface, all 32-bit write activity is seen at the bus interface as single write transactions from the CPU core. There is no such thing as a “four word block burst write” from the CPU core; the processor performs a word or sub-word write as a single autonomous bus transaction. However, the on-chip DMA channels are capable of generating multi-word burst writes if so programmed. The `SysBurstFrame` output is used to decode burst writes. Successive write transactions can be processed using page mode writes by DRAM Memory Controller. This is particularly important when “flushing” the write buffer before performing a data read.

Uncached writes—which contain only 1, 2, or 3 bytes of data—assert the appropriate byte enables, `MemWrEn(3:0)` or `DramCAS(3:0)`. Thus, there really is only one type of non-burst 32-bit write transaction. However, in some cases such as with the DRAM Controller, the memory system may elect to take advantage of the assertion of a page comparator internal write near signal during a write to perform quicker write operations than would otherwise be performed.

In processing 32-bit writes, there is only one parameter of interest: the latency of the write. This latency is influenced by the overall system architecture as well as the type of memory system being addressed: time required to perform address decoding and bus arbitration, memory pre-charge requirements, and memory control requirements, as well as memory access time.

The R36100 has been designed to accommodate a wide variety of memory system designs, including no-wait cycle operations (write completed in two cycles) through simpler, slower systems incorporating many bus wait cycles.

16-Bit Transactions

When the R36100 uses a 16-bit port, it does its writes in halfword size or smaller increments. Thus if the data contains 8 or 16 bits (1 or 2 bytes), it will be handled with a single halfword write with the appropriate byte enables, `MemWrEn(1:0)` or `DramCAS(1:0)` asserted. Note that during 16-bit accesses, `MemWrEn` and `DramCAS` bit3 is equal to bit 0 and bit 2 is equal to bit 1.

If the data contains 24 or 32 bits (3 or 4 bytes), it will be handled with a double halfword write mini-burst with the appropriate byte enables, `MemWrEn(1:0)` or `DramCAS` for each halfword asserted. A mini-burst puts both halfwords out as separate data phases of the same write transaction. The memory system simply returns an internal AckN for each halfword datum which will automatically increment `SysAddr(3:1)` and change the write enables if appropriate.

Similar to a read mini-burst, a write mini-burst can be detected using the `SysBurstFrame` signal to determine when the final halfword data is being returned or by using the de-assertion of the `SysWr` line. The R36100 is designed to accommodate a wide variety of different memory bandwidths, including DRAM systems that need pre-charge wait cycles for the first halfword and then use a fast page mode access for bursting the second halfword.

The data lines used in 16-bit ports are always SysData(31:16) for big endian systems and SysData(15:0) for little endian systems. This is regardless of the Reverse Endianess bit in the CP0 Status register. For big endian systems, $\overline{\text{MemWrEn}}(3)$ and $\overline{\text{DramCAS}}(3)$ correspond to the byte lane in SysData(31:24) and $\overline{\text{MemWrEn}}(2)$ and $\overline{\text{DramCAS}}(2)$ correspond to SysData(23:16). For little endian systems, $\overline{\text{MemWrEn}}(1)$ and $\overline{\text{DramCAS}}(1)$ correspond to the byte lane in SysData(15:8) and $\overline{\text{MemWrEn}}(0)$ and $\overline{\text{DramCAS}}(0)$ correspond to SysData(7:0).

8-Bit Transactions

When the R36100 uses an 8-bit port, it performs writes in byte size increments. Thus if the data contains 1 byte, it will be handled with a single byte write. If the data contains 2, 3, or 4 bytes, it will be handled with a double, tri, or quad byte write mini-burst, respectively. A mini-burst puts 2, 3, or 4 bytes out as separate data phases of the same write transaction.

The memory system simply returns an internal AckN for each byte datum which will automatically increment SysAddr(3:0). Similar to a read mini-burst, a write mini-burst can be detected using the $\overline{\text{SysBurstFrame}}$ signal to determine when the final byte datum is being returned or by using the de-assertion of the $\overline{\text{SysWr}}$ line. The R36100 is designed to accommodate a wide variety of different memory bandwidths, including DRAM systems that need pre-charge wait cycles for the first byte and then use a fast page mode access for bursting subsequent bytes.

The data lines used in 8-bit ports are always SysData(31:24) for big endian systems and SysData(7:0) for little endian systems. This is regardless of the Reverse Endianess bit in the CP0 Status register.

Write Interface Timing Overview

The protocol for transmitting data from the processor to memory and I/O devices are discussed below.

Initiating the Write

A write transaction occurs when the processor has placed data into the write buffer, and the bus interface is either free, or write has the highest priority. Internally, the processor bus arbiter uses the NotEmpty indicator from the write buffer to indicate that a write is being requested.

Assuming that the write transaction can be processed (that is, there are no ongoing bus operations, and no higher priority operations pending), the processor will initiate a bus write transaction on the next rising edge of SysClk. Higher priority operations would have the effect of delaying the start of the write.

Figure 7.8 on page 21 illustrates the initiation of a write transaction, based on the internal negation of the internal WbEmptyN control signal. This figure applies when the processor is performing a write, and the write buffer is otherwise empty. If the write buffer already had data in it, the buffer would continually request the use of the bus until it was emptied; it would be up to the bus interface unit arbiter to decide the priority of the request relative to other pending requests. Additional stores would be captured by other write buffer entries, until the write buffer was filled.

Memory Addressing

A write transaction begins when the processor asserts its $\overline{\text{SysWr}}$ control output, and also drives the address and other control information onto the SysAddr and memory interface buses. The data is driven with $\overline{\text{SysALEn}}$ asserting. Figure 7.8 also illustrates the start of this type of processor write transaction, including the addressing of memory and presenting the store data on the SysData bus.

At the rising edge of $\overline{\text{SysClk}}$, the processor will drive the write target address onto the SysAddr bus. At this time, $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurst-Frame}}$ will also be asserted, to indicate to external ASICs and peripherals that a memory transaction is beginning.

The Data Phase

Simultaneous with driving the address out, the data phase begins.

The processor enters the data phase by performing the following sequence of events:

- It negates $\overline{\text{SysALEn}}$.
- It internally captures the data in a register in the bus interface unit, and enables this register onto its output drivers on the SysData bus.

At this time, it begins to look for the end of the write cycle.

Terminating the Write

There are only two methods for the external memory system to terminate a write operation:

- It can supply the appropriate number of internal AckNs (acknowledges) to the processor by using an internal memory controller wait-state generator to indicate that it has sufficiently processed the write request, and that the processor may terminate the write.

It can supply a $\overline{\text{SysBusError}}$ to the processor, to indicate that the requested data transfer has “failed” on the bus. The system interface behavior of the processor when $\overline{\text{SysBusError}}$ is presented is identical to the behavior when the last internal AckN is asserted. In the case of writes terminated by $\overline{\text{SysBusError}}$, no exception is taken, and the data transfer cannot be retried.

Figure 7.9 shows the timing of the control signals when the write cycle is being terminated.

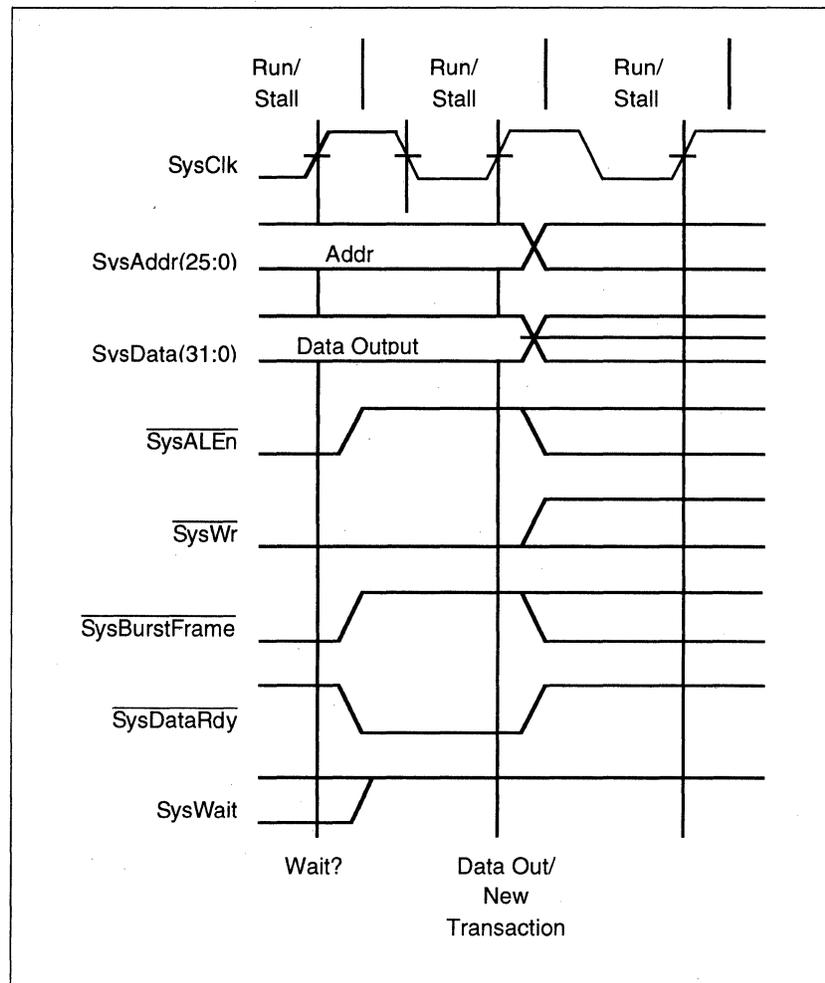


Figure 7.8 Write Cycle Termination

To determine the end of the write cycle one of these methods may be used:

- Systems that only use 32-bit memory sub-region ports, such as the R3051 family, only have single datum writes as generated from the CPU and either count the number of wait-cycles or use the deasserting edge of $\overline{\text{SysWr}}$ to end the transaction. However, since the on-chip DMA Controller can generate burst writes, memory systems in general must be able to handle bursts.
- Systems that use 16 or 8-bit ports must in general support mini-burst writes. Memory controllers for such systems can use the deasserting edge of $\overline{\text{SysWr}}$ to reset the controller. An external memory controller or logic analyzer can also look for $\overline{\text{SysBurstFrame}}$ to deassert. When $\overline{\text{SysBurstFrame}}$ deasserts, the controller knows that it is handling the final datum of the transaction.

Latency Between Processor Operations

In general, the processor may begin a new bus activity in the phase immediately after the termination of the write cycle. This operation may be either a read, write, or bus grant. A new operation may begin as soon as one clock cycle after the final internal AckN is sampled from the interface.

Also note that bus turn around after a write transaction does not occur. That is, the processor continues to drive the SysAddr and SysData buses throughout the write transaction (both address and data phases), and it will also drive the SysAddr bus during the start of either a subsequent read or write transaction. Since no change in bus ownership occurs, the Bus Turn Around field of the CPO Bus Control register does not apply after write transactions.

Write Buffer Full Operation

It is possible that the execution core on occasion may be able to fill the on-chip write buffer. If the processor core attempts to perform a store to the write buffer while the buffer is full, the execution core will be stalled by the write buffer until a space is available. Once space is made available, the execution core will use an internal fixup cycle to “retry” the store, allowing the data to be captured by the write buffer. It will then resume execution.

The write buffer can actually be thought of as “four and one-half” entries: it contains a special data buffer which captures the data being presented by an ongoing bus write transaction. Thus, when the bus interface unit begins a write transaction, the write buffer slot containing the data for that write is freed up in the second phase of the write transaction. If the processor was in a write busy stall, it will be released to write into the now available slot at this time, regardless of how long it takes the memory system to retire the ongoing write.

Note that each write buffer entry is one internal 32-bit word wide, but each entry can only hold the result of one store operation. Thus a 32-bit port can store 4 words while a 16-bit port can store up to 8 halfwords when using store word operands. However, if for example, four store byte operations are done, each byte takes a full entry.

Figure 7.9 illustrates the write-buffer-full operation.

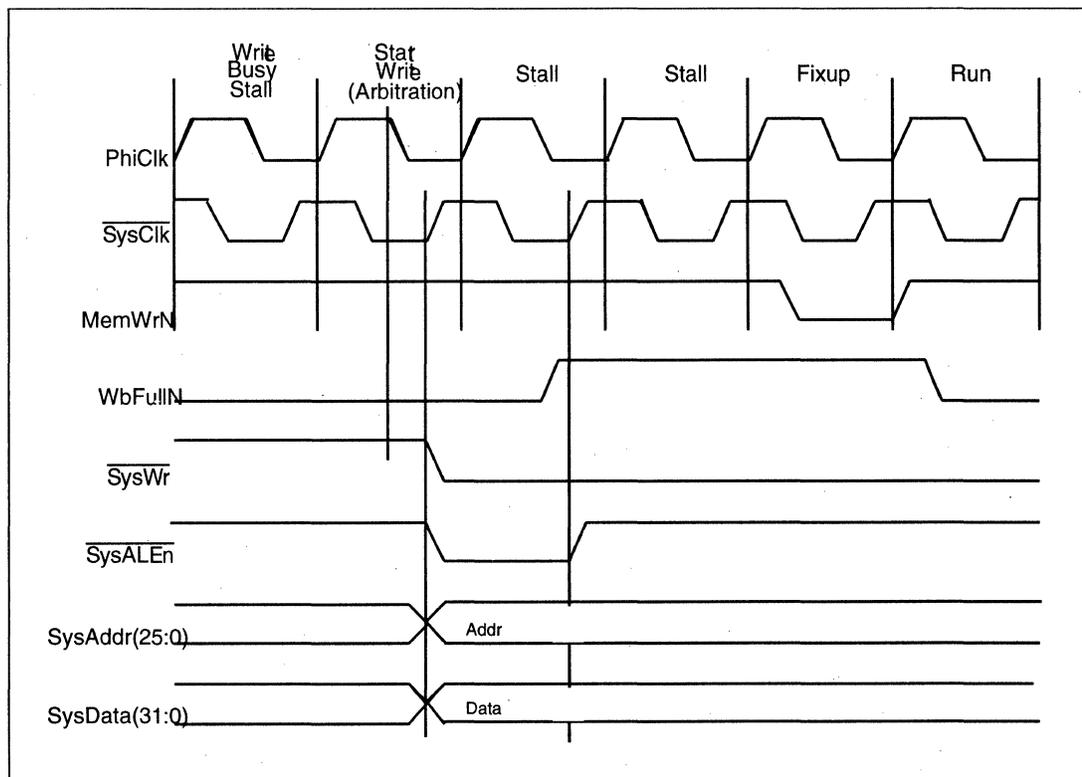


Figure 7.9 Write-Buffer-Full Operation



Integrated Device Technology, Inc.

Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. One of the on-chip bus controllers is the **"Memory Controller"** as described in this chapter.

This chapter will provide an overview of the Memory Controller interface, a complete description of the signal pins and their timing, and how the interface relates to typical external hardware ROMs and RAMs.

Features

- Controls ROM, Flash, EEPROM, SRAM and PCMCIA style memories
- Controls up to 8 banks of memory
(Note: chip selects are shared with the I/O Controller)
- Interleaved and Non-Interleaved support
- Each MemCS can be programmed to:
 - Individual chip selects
 - Combined interleaved pair-wise chip selects
 - Combined PCMCIA/MEM-style pair-wise chip selects
- Each Bank has Programmable Base Address
- Each Bank Size programmable from 8KB - 64MB
- 8, 16, 32-bit, and interleaved 32-bit boot prom support
- Wait State Generator features:
 - Programmable time from start to end of each data access for each area
 - Programmable time options for reads and writes
 - Programmable time options for single word accesses and for burst block accesses
 - Internally generates the RdCEnN and AckN timing for all CPU accesses
 - A programmed value may be overridden by the $\overline{\text{SysWait}}$ input signal
 - Direct control of data path transceivers supports various options:
 - Direct Bus Connection
 - FCT260 Bidirectional Latched Multiplexor
 - FCT245 Bidirectional Transceiver
 - FCT543 Bidirectional Registered Transceiver

Block Diagram

The functional block diagram of the Memory Bus Controller is shown in Figure 8.1. Starting at the bottom, the main Memory Controller Control Signal State Machine is responsible for generating the basic Memory Control signals used to connect to external PROMs, SRAMs, and other similar types of memory. These signals include chip selects, read enables/strobes, and write enables/strobes. The Memory Controller as a whole works in cooperation with the Bus Interface Unit described in Chapter 7.

Thus the Control Signal State Machine sends and receives information from the BIU Controller for assistance with controlling the port width and controlling partial word reads and writes. The Control Signal State Machine also uses information stored in the software programmable Memory Controller Register Bank for example, to control interleaved versus non-interleaved memory cycles.

The Memory Controller Wait-State Generator is located in the center of Figure 8.1. The Wait-State Generator takes care of sending and receiving information from the BIU Controller in order to control the sequencing and timing of reading and writing each individual datum. The number of wait-states is derived from the settings programmed into the Register Bank. Once the correct number of wait-states has been counted out, then the Wait-State Generator sets the appropriate internal BIU Acknowledge signals. With the programmable Wait-State Generator it is possible to eliminate the external state machines that are traditionally used for this function.

At the top of Figure 8.1 is the Memory Controller Decoder. The decoder constantly monitors the Bus Interface Unit's address and data bus to see if either:

1. The access is to the Memory Controller's Register Bank.
2. The access is in one of the Memory Controller's Chip Select Areas that is responsible for controlling the bus transaction.

And, at the left of Figure 8.1, is the Memory Controller Register Bank. The Register Bank allows the software programmer access to the many different options of the Memory Controller.

The chip select address ranges, the number of wait-states, the port-width of the chip select, and other similar options are programmed into the Register Bank as part of the software initialization sequence of the boot operating system.

Because the Memory Controller is typically used by the boot PROM, the essential default values of the boot PROM chip select, MemCS(0) are set by the Reset Initialization Vector as described in Chapter 18, "Reset and Clocking."

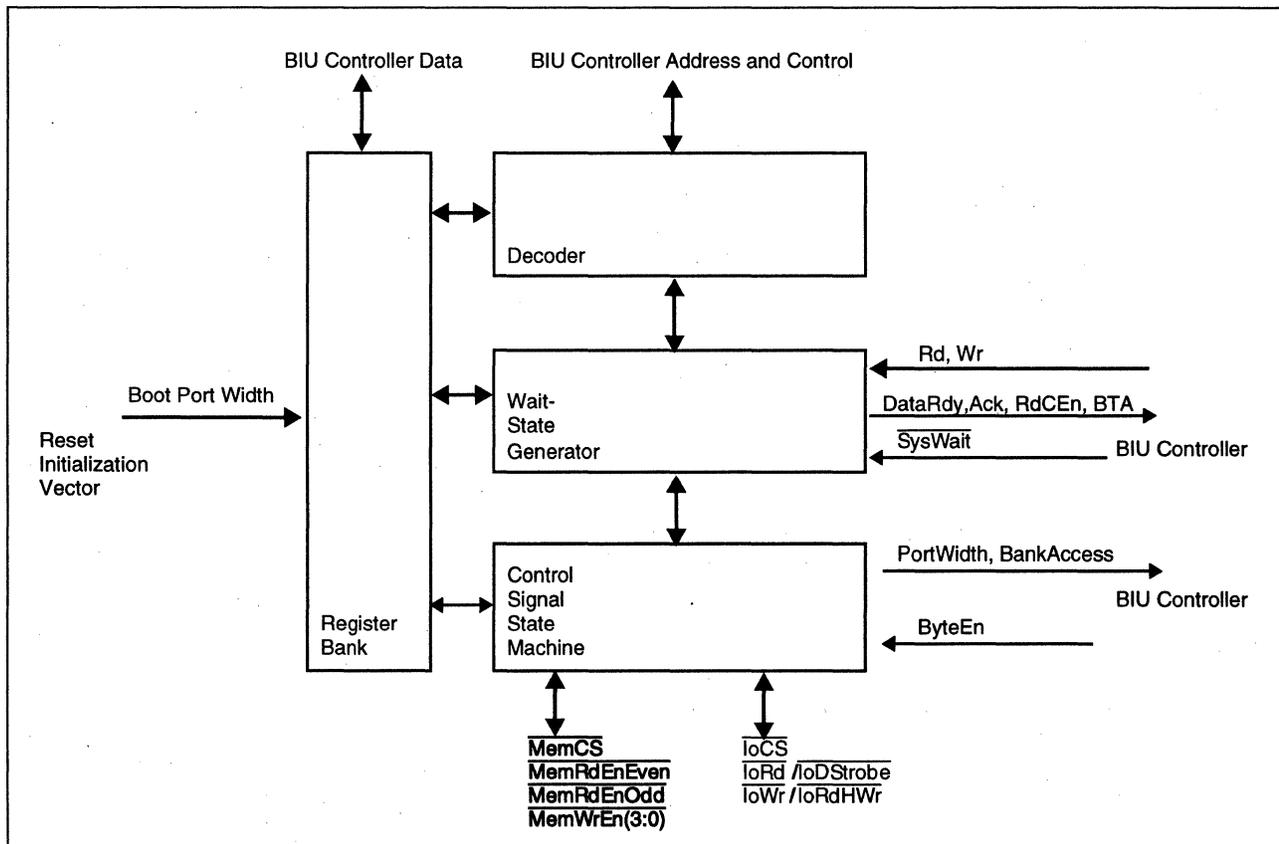


Figure 8.1 R36100 Memory Bus Controller Block Diagram

Memory Bus Controller Interface Signals

Memory Controller Signals

These external pins are typically attached directly from the R36100 RISController to external ROM and RAM chips and their transceivers.

MemCS(7:0) / **Output** IoCS(7:0)

Memory Chip Select: The $\overline{\text{MemCS}}$ signals are active low outputs used to select one of the programmable memory controller areas. Typically a bank of external memory chips each attach to a $\overline{\text{MemCS}}$ signal such that the memory bank can be selected and turned on during a memory transaction. When the address from the CPU or DMA Controller matches the memory block corresponding to a particular $\overline{\text{MemCS}}$ signal, that $\overline{\text{MemCS}}$ will assert at the beginning of the next memory transaction and de-asserts at the end of that transaction.

$\overline{\text{MemCS}}$ signals are used individually for non-interleaved systems or in pairs (i.e., $\overline{\text{MemCS}}$ 0 & 1, 2 & 3, 4 & 5, or 6 & 7) for interleaved systems. When using interleaved memory, the pair of chip selects are both asserted for a read but only one at a time is asserted for a write.

The boot PROM is assigned to $\overline{\text{MemCS}}(0)$ and if interleaved, $\overline{\text{MemCS}}(1)$. The port width option is determined using the Reset Initialization Vector on the $\overline{\text{ExcInt}}(2:0)$ pins. Other options are set to universal settings which the boot software can reprogram.

The $\overline{\text{MemCS}}$ chip selects are selectable and shared between Memory and I/O type-types (see Chapter 10, "I/O Bus Controller").

MemRdEnEven **Output**

Memory Read Enable Even Bank: This active-low output signal is used as a read enable strobe used in conjunction with the even chip selects, such as $\overline{\text{MemCS}}(6:4:2:0)$. Typically, $\overline{\text{MemRdEnEven}}$ is attached to all even bank's memory chips and transceivers (if present). If the banks are interleaved, this signal is the output enable for the even bank of the selected memory pair. Whether the banks are interleaved or non-interleaved, all banks that share the same transceiver must use either all even or all odd chip selects ($\overline{\text{MemCSs}}$) rather than an odd/even pair unless external gating circuitry is provided. $\overline{\text{MemRdEnEven}}$ controls when the memory chip and its transceiver (if present) can drive the data signals back on to the main system data bus, SysData(31:0).

Transceivers (FCT260, FCT245, or FCT543) can be used to isolate memory systems and allow for the different turnoff times for various memory devices. If transceivers are used, then during Multiplexer FCT260-Type interleaved accesses, $\overline{\text{MemRdEnEven}}$ OR's with $\overline{\text{MemRdEnOdd}}$ internally such that $\overline{\text{MemRdEnEven}}$ remains asserted for the majority of the bus transaction for both even and odd accesses. Typically in the Multiplexer Mode, $\overline{\text{MemRdEnEven}}$ is attached to the common output enable input pin of the multiplexer while $\overline{\text{MemRdEnOdd}}$ is attached to the path input pin of the multiplexer. During FCT245-Type interleaved type accesses, $\overline{\text{MemRdEnEven}}$ OR's with $\overline{\text{MemWrEnEven}}$ internally.

MemRdEnOdd**Output**

Memory Read Enable Odd Bank: This active-low output signal is used as a read enable strobe used in conjunction with the odd chip selects, (i.e., $\overline{\text{MemCS}}(7:5:3:1)$). Typically $\overline{\text{MemRdEnOdd}}$ is attached to all odd bank's memory chips and transceivers (if present). During FCT245-Type interleaved type accesses, $\overline{\text{MemRdEnOdd}}$ OR's with $\overline{\text{MemWrEnOdd}}$ internally. Please see the signal description for $\overline{\text{MemRdEnEven}}$ for more information.

MemWrEn(3:0) /**Output/(Input during DMA)****MemByteEn(3:0)****MemAddr(29:26) /**

System Write Enable: This dedicated byte enable strobes is always driven by the R36100, except during the Address Strobe of external DMA cycles. The pins always act as write enables except during PCMCIA-type accesses or PCI-style accesses.

Output - During 32-bit accesses, these strobes are always de-asserted, except for the appropriate byte lane(s) for partial-word writes and mini-bursts. SRAM and Flash EPROM memories can directly connect their byte write enables to the R36100 $\overline{\text{MemWrEn}}(3:0)$ signals. During 16-bit accesses, either $\overline{\text{MemWrEn}}(3:2)$ or $\overline{\text{MemWrEn}}(1:0)$ are used as both pairs will be equivalent. During 8-bit accesses, either $\overline{\text{MemWrEn}}(3)$ or $\overline{\text{MemWrEn}}(0)$ are used as they will be equivalent.

Input - During External DMA accesses, the external DMA Controller must assert the appropriate byte lanes during Phase 1 of a write when $\overline{\text{SysALEn}}$ and $\overline{\text{SysWr}}$ are asserted. The byte enables are ignored by the BIU Controller during external DMA reads.

$\overline{\text{MemWrEn}}(3:0)$ assert and de-assert on the falling edge of $\overline{\text{SysClk}}$, whereas most control signals use the opposite edge of the clock. Thus it is not generally recommended that they be used for any external state machine inputs.

During idle cycles, $\overline{\text{MemWrEn}}(3:0)$ will return to inactive (high).

Memory Byte Enable Bus: During a PCI-Memory-style or PCI-I/O-style access, the $\overline{\text{MemWrEn}}(3:0)$ bus is programmed to assert the appropriate byte lanes on both reads and writes, rather than just on writes. In this mode as with the other modes, $\overline{\text{MemByteEn}}(3:0)$ are required to return de-asserted high at the end of every bus transaction.

Memory Address Bus: During a PCMCIA Memory or I/O type access, the $\overline{\text{MemWrEn}}(3:0)$ bus is instead driven with Physical Address bits. On the R36100 and R3051-base family memory map, virtual and physical addresses (29:0) are the same. An application using PCMCIA can for example use $\overline{\text{MemAddr}}(27:26)$ to externally decode PCMCIA style chip selects into as many as four ($256\text{M}/64\text{M} = 4$) slots. In this mode as in the other modes, the signals all return inactive high at the end of the bus transaction.

BIU Controller Signals

Many of the BIU Controller Signals are necessary to complete the memory interface. These signals are listed here as a reminder. Information specific to the Memory Controller is given here and general information about the signal is given in Chapter 7, "BIU Controller."

SysAddr(25:0) Output/Input

System Address Bus: SysAddr is an output bus when used with the Memory Controller. Typical 32-bit memory banks connect the word offset of the Least Significant Bits (LSB) of SysAddr to each memory chip. Thus the typical 32-bit memory bank skips SysAddr(1:0) and connects starting with SysAddr(2) on up. Typical 16-bit memory banks connect the half-word offset of the LSBs of SysAddr to each memory chip starting with SysAddr(1) on up. Typical 8-bit memory banks connect the LSBs of SysAddr to each memory chip starting with SysAddr(0) on up.

SysData(31:0) Output/Input

System Data Bus: Typical 32-bit memory banks connect the entire 32-bit SysData bus to the memory chips' data pins or to their transceivers. *16-bit and 8-bit memory banks connect to particular data pins depending on whether the Endianness of the system is Big Endian or Little Endian.* Thus 16-bit memory banks use SysData(31:16) if they are Big Endian and SysData(15:0) if they are Little Endian. 8-bit memory banks use SysData(31:24) if they are Big Endian and SysData(7:0) if they are Little Endian. The User Mode Reverse Endianness Bit in the CP0 Status Register has no effect on the connections to SysData, however, it is strongly recommended that the Reverse Endianness Bit not be used to "correct" an endianness connection as it does not function in the kernel mode, nor does it correct the boot prom endianness.

SysWait Input

System Wait Negated: $\overline{\text{SysWait}}$ can be used by an external source to add wait-states to the Memory Controller. Since the Memory Controller itself has a Wait-State Generator, $\overline{\text{SysWait}}$ typically is not needed and can be pulled-up with a resistor. The most likely application of $\overline{\text{SysWait}}$ is for an asynchronous memory event such as a Dual-Port Memory Busy signal which can be used to attach to $\overline{\text{SysWait}}$ to delay the beginning of a Memory Transaction in the Wait Mode option of the Memory Control Register 2. Please see Chapter 7. "BIU Controller" for a general description of $\overline{\text{SysWait}}$.

Overview of the Memory Controller (MemCntrl)

The Memory Controller (MemCntrl) provides control for all memory spaces except for the DRAM space. These memory areas are intended for use by items such as boot ROM, Flash memory cards, additional EPROM and EEPROM space, SRAMs, and Dual Port RAM. Such memories typically have address inputs, data I/O, chip select, read output enable, and if writable, a write enable strobe.

Chip Selects

The Memory Controller (MemCntrl) contains up to 8 separate memory spaces each having its own Memory Controller Chip Select ($\overline{\text{MemCS}}$) output pin. Each $\overline{\text{MemCS}}$ space occupies from 8K to 256MB of address space of which 64MB is externally addressable (due to the 26 address lines). The address space that each $\overline{\text{MemCS}}$ decodes is programmable. The MemCntrl will use the programmed information in the MSB and LSB Base Address Registers along with the size (8K to 256MB) of the given area as programmed in the MSB and LSB Page Mask Registers. This information is used to compare with the address asserted by the CPU-BIU or DMA Controller to determine if that particular $\overline{\text{MemCS}}$ area is being accessed for the current read or write. Each area supports single reads, burst reads, single writes, and burst writes. The port size of the data path (8, 16, 32-bit, or interleaved) of each area is also programmable with each area's Control Register.

The $\overline{\text{MemCS}}$ signals can be used in pairs for interleaving. The pairs are $\overline{\text{MemCS}}(1:0)$ for one interleaved area and for the others, $\overline{\text{MemCS}}(3:2)$, $\overline{\text{MemCS}}(5:4)$, and $\overline{\text{MemCS}}(7:6)$. When interleaved, both chip selects within a pair must be programmed to the same values by the user. Note that the memory controller does not support seamless jumperless upgrades from single bank to interleaved systems because of the system dependent address multiplexing involved. This can be done externally for RAM types, however, for ROM types, the PROM chip programmer would require "switching of the binary object file" which is rarely worth the trouble.

Transceiver Control Interface

The Memory Controller provides transceiver output enables and write enables that are suitable for direct bus connection, FCT260 multiplexors, FCT245 transceivers, or for FCT543 bidirectional registers. The selection of the type of memory is software programmable. 8/16/32-bit wide Boot PROMs can use either the direct bus connection, FCT244s, or FCT543s. FCT245s can be used if the Boot PROM initializes the Memory Controller before doing any writes. Interleaved Boot PROMs are assumed to use the FCT260 type.

Wait State Generator

The Wait-State Generator (WSG) controls the speed of the memory accesses to and from the Bus Interface Unit Controller. This includes the time from the start of a memory transaction until the first datum is sent or received and the time between consecutive datum on burst transactions. The WSG also is programmed to generate the internal RdCEnN and AckN signals for CPU read and write requests.

The internal Acknowledge signal, "AckN" (described in Chapter 7) is the same as the external signal pin that the R3051 RISController family uses. On single word reads and on both single word and burst writes, AckN is automatically placed at the end of the transaction by the WSG. However, on burst reads, the user is required to program the correct value into each corresponding MemCS area's Control Register "Burst Ack" field in order to optimize the CPU pipeline restart after a burst read. The most optimal value is 3 pipeline clocks previous to when the last datum is sampled. Less optimal values can be used if for instance SysWait is used.

The signal called SysWait can be used to override the programmed settings of the Wait State Generator. The actual action the WSG performs when SysWait is asserted will depend on when it is asserted relative to the transaction. SysWait has a pipeline delay, such that it must be asserted two clocks before the desired effect is noticeable. By asserting it immediately after a datum is received or transmitted, the next datum can be delayed. However, use for this purpose is generally not recommended since the WSG has the same functionality.

The SysWait signal is useful for accessing memories such as Dual-Port-type memory, off-card "Ack"-type memory, and PCI-style memory.

Register Option Programmability

The Memory Controller contains 8 sets of registers, one set for each chip select. These registers allow the Memory Controller to be configured for different types and speeds of memory. Thus almost any system speed/cost/manufacturing trade-off can be accommodated.

Register Descriptions

The Memory Controller Registers are divided into 8 sets of registers, one set for each chip select memory area. Table 8.1 and Table 8.2 provide the address map and descriptions of the Memory and I/O Controller registers. These registers are shared with the I/O Controller as described in Chapter 9.

Phys. Address	Register Description
0xFFFF_E200	Memory and I/O LSB Base Address Register for Bank 0
0xFFFF_E204	Memory and I/O MSB Base Address Register for Bank 0
0xFFFF_E208	Memory and I/O LSB Bank Mask Register for Bank 0
0xFFFF_E20C	Memory and I/O MSB Bank Mask Register for Bank 0
0xFFFF_E210	Memory and I/O Control Register for Bank 0
0xFFFF_E214	Memory and I/O LSB Wait State Generator Register for Bank 0
0xFFFF_E218	Memory and I/O MSB Wait State Generator Register for Bank 0
0xFFFF_E220	Memory and I/O LSB Base Address Register for Bank 1
0xFFFF_E224	Memory and I/O MSB Base Address Register for Bank 1
0xFFFF_E228	Memory and I/O LSB Bank Mask Register for Bank 1
0xFFFF_E22C	Memory and I/O MSB Bank Mask Register for Bank 1
0xFFFF_E230	Memory and I/O Control Register for Bank 1
0xFFFF_E234	Memory and I/O LSB Wait State Generator Register for Bank 1
0xFFFF_E238	Memory and I/O MSB Wait State Generator Register for Bank 1
0xFFFF_E240	Memory and I/O LSB Base Address Register for Bank 2
0xFFFF_E244	Memory and I/O MSB Base Address Register for Bank 2
0xFFFF_E248	Memory and I/O LSB Bank Mask Register for Bank 2
0xFFFF_E24C	Memory and I/O MSB Bank Mask Register for Bank 2
0xFFFF_E250	Memory and I/O Control Register for Bank 2
0xFFFF_E254	Memory and I/O LSB Wait State Generator Register for Bank 2
0xFFFF_E258	Memory and I/O MSB Wait State Generator Register for Bank 2
0xFFFF_E260	Memory and I/O LSB Base Address Register for Bank 3
0xFFFF_E264	Memory and I/O MSB Base Address Register for Bank 3
0xFFFF_E268	Memory and I/O LSB Bank Mask Register for Bank 3
0xFFFF_E26C	Memory and I/O MSB Bank Mask Register for Bank 3
0xFFFF_E270	Memory and I/O Control Register for Bank 3
0xFFFF_E274	Memory and I/O LSB Wait State Generator Register for Bank 3
0xFFFF_E278	Memory and I/O MSB Wait State Generator Register for Bank 3
0xFFFF_E280	Memory and I/O LSB Base Address Register for Bank 4
0xFFFF_E284	Memory and I/O MSB Base Address Register for Bank 4
0xFFFF_E288	Memory and I/O LSB Bank Mask Register for Bank 4
0xFFFF_E28C	Memory and I/O MSB Bank Mask Register for Bank 4
0xFFFF_E290	Memory and I/O Control Register for Bank 4
0xFFFF_E294	Memory and I/O LSB Wait State Generator Register for Bank 4
0xFFFF_E298	Memory and I/O MSB Wait State Generator Register for Bank 4
NOTES: 1. Big Endian software must offset these addresses by b'10 (0x2).	

Table 8.1 List of the Memory and I/O Controller Registers (1 of 2).

Phys. Address	Register Description
0xFFFF_E2A0	Memory and I/O LSB Base Address Register for Bank 5
0xFFFF_E2A4	Memory and I/O MSB Base Address Register for Bank 5
0xFFFF_E2A8	Memory and I/O LSB Bank Mask Register for Bank 5
0xFFFF_E2AC	Memory and I/O MSB Bank Mask Register for Bank 5
0xFFFF_E2B0	Memory and I/O Control Register for Bank 5
0xFFFF_E2B4	Memory and I/O LSB Wait State Generator Register for Bank 5
0xFFFF_E2B8	Memory and I/O MSB Wait State Generator Register for Bank 5
0xFFFF_E2C0	Memory and I/O LSB Base Address Register for Bank 6
0xFFFF_E2C4	Memory and I/O MSB Base Address Register for Bank 6
0xFFFF_E2C8	Memory and I/O LSB Bank Mask Register for Bank 6
0xFFFF_E2CC	Memory and I/O MSB Bank Mask Register for Bank 6
0xFFFF_E2D0	Memory and I/O Control Register for Bank 6
0xFFFF_E2D4	Memory and I/O LSB Wait State Generator Register for Bank 6
0xFFFF_E2D8	Memory and I/O MSB Wait State Generator Register for Bank 6
0xFFFF_E2E0	Memory and I/O LSB Base Address Register for Bank 7
0xFFFF_E2E4	Memory and I/O MSB Base Address Register for Bank 7
0xFFFF_E2E8	Memory and I/O LSB Bank Mask Register for Bank 7
0xFFFF_E2EC	Memory and I/O MSB Bank Mask Register for Bank 7
0xFFFF_E2F0	Memory and I/O Control Register for Bank 7
0xFFFF_E2F4	Memory and I/O LSB Wait State Generator Register for Bank 7
0xFFFF_E2F8	Memory and I/O MSB Wait State Generator Register for Bank 7

NOTES: 1. Big Endian software must offset these addresses by b'10 (0x2).

Table 8.2 List of the memory and I/O Controller Registers (2 of 2)

Memory MSB Base Address Register for Bank 7..0
 ('MemMSBBaseAddrReg(7..0)'),
 and
Memory LSB Base Address Register for Bank 7..0
 ('MemLSBBaseAddrReg(7..0)')

There are 8 pairs of Base Address MSB & LSB Registers, a pair for each Memory Chip Select (MemCS). Each pair of memory base address registers is concatenated into an internal "32-bit" register and refer to the most significant 16 address bits and the least significant 16 address bits.

The formats of the MemMSBBaseAddrReg and MemLSBBaseAddrReg are displayed in Figure 8.2 and Figure 8.3. These registers are both readable and writable.

The Base Address Registers are used to determine the starting location of a particular Memory Chip Select.

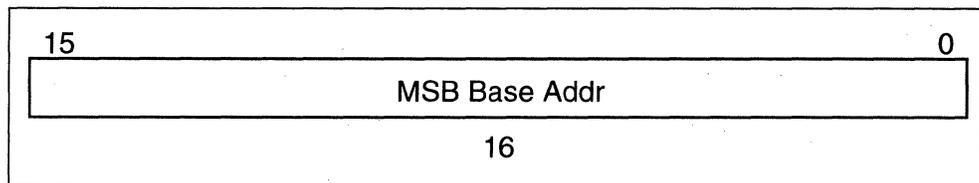


Figure 8.2 Memory and I/O MSB Base Address Register ('MemMSBBaseAddrReg').

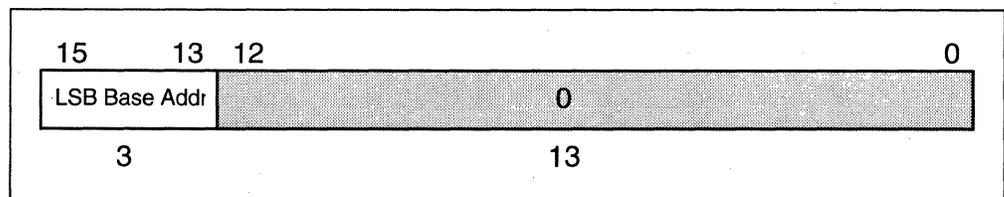


Figure 8.3 Memory and I/O LSB Base Address Register ('MemLSBBaseAddrReg').

Because of the possibility of interleaving, there are 4 groups of 2 chip selects, as follows:

Group 0:	$\overline{\text{MemCS}}(1:0)$
Group 1:	$\overline{\text{MemCS}}(3:2)$
Group 2:	$\overline{\text{MemCS}}(5:4)$
Group 3:	$\overline{\text{MemCS}}(7:6)$

Bits 31-28 of each group must be programmed identically since the internal hardware uses bits 31-28 from the even register, $\overline{\text{MemCS}}(0,2,4,6)$ for each group. This corresponds to setting each group of four chip selects into 1 of 16 possible 256M address spaces.

Bits 27-13 must be programmed to the desired base address. This corresponds to separate address spaces for each chip select of 8K to 256M.

Internally, bits 12:0 are reserved to '0' and must be programmed as '0'. This corresponds to having minimum contiguous chip select banks of 8K. The default base addresses at reset are shown in Table 8.3.

Chip Select	Default Value
$\overline{\text{MemCS}}(0)$	0x 1FC0_0000
$\overline{\text{MemCS}}(1)$	0x 1FF0_0000 non-interleaved (if interleaved, then 0x1FC00000)
$\overline{\text{MemCS}}(2)$	0x 2FC0_0000
$\overline{\text{MemCS}}(3)$	0x 2FF0_0000 non-interleaved (if interleaved, then 0x 2FC0_0000)
$\overline{\text{MemCS}}(4)$	0x 3FC0_0000
$\overline{\text{MemCS}}(5)$	0x 3FF0_0000 non-interleaved (if interleaved, then 0x 3FC0_0000)
$\overline{\text{MemCS}}(6)$	0x 4FC0_0000
$\overline{\text{MemCS}}(7)$	0x 4FF0_0000 non-interleaved (if interleaved, then 0x 4FC0_0000)

Table 8.3 Memory and I/O Controller Base Addresses.

In summary:

1. Bits 31:28 of each group of four $\overline{\text{MemCS}}$ base addresses is set by the first $\overline{\text{MemCS}}$ of the group.
2. Bits 27:13 of each $\overline{\text{MemCS}}$ base address is used to distinguish the starting address of each memory space within a group.
3. Bits 12:0 are always ignored.

Memory MSB Bank Mask Register for Bank 7..0
('MemMSBBankMaskReg(7..0)'),
and

Memory LSB Bank Mask Register for Bank 7..0
('MemLSBBankMaskReg(7..0)')

There are 8 pairs of Bank Mask Registers, one pair for each chip select ($\overline{\text{MemCS}}$). The two Bank Mask Registers are concatenated into an internal "32-bit" register and refer to the most significant 16 address bits and the least significant 16 address bits.

The formats of the MemMSBBankMaskReg and MemLSBBankMaskReg are displayed in Figure 8.4 and Figure 8.5. These registers are both readable and writable and bits 31:13 are set and bits 12:0 are cleared by default on reset.

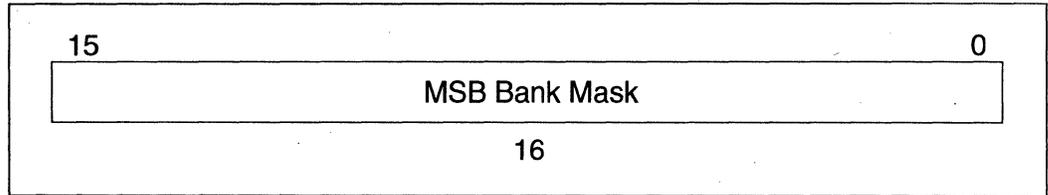


Figure 8.4 Memory and I/O MSB Page Mask Register ('MemMSBPageMaskReg').

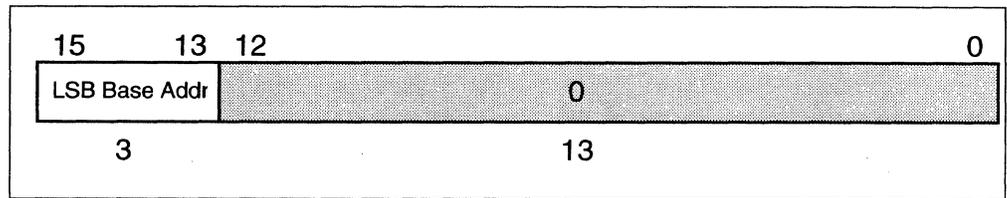


Figure 8.5 Memory and I/O LSB Page Mask Address Register ('MemLSBPageMaskReg').

The Bank Mask Registers are used to decide which address bits in the base address are to be used for comparing whether a chip select ($\overline{\text{MemCS}}$) is to be activated.

The internal grouping of chip selects is as follows:

- Group 0: $\overline{\text{MemCS}}(1:0)$
- Group 1: $\overline{\text{MemCS}}(3:2)$
- Group 2: $\overline{\text{MemCS}}(5:4)$
- Group 3: $\overline{\text{MemCS}}(7:6)$

Bits 31-28 of each group must be programmed identically since the internal hardware uses bits 31-28 from the even register, $\overline{\text{MemCS}}(0,2,4,6)$ for each group. This corresponds to setting each group of four chip selects into 1 of 16 possible 256M address spaces.

Internally, bits 27-13 must be programmed to the desired Bank mask. This corresponds to separate address spaces for each chip select of 8K to 256M. Bits 12:0 are reserved to '0' and must be programmed as '0'. This corresponds to having minimum contiguous chip select banks of 8K.

In summary:

1. Bits 31:28 of each group of four $\overline{\text{MemCS}}$ Bank masks is set by the first $\overline{\text{MemCS}}$ of the group.
2. Bits 27:13 of each $\overline{\text{MemCS}}$ Bank mask are used to distinguish the size of each memory space.
3. Bits 12:0 are always ignored

Table 8.4 lists the values and actions for the Memory Mask Field Encoding, Figure 8.6 shows the Memory and I/O Control Register, and Table 8.5 provides the register's bit assignments.

Value	Action
'1'	Bit is used in Address comparison
'0'	Bit is masked out

Table 8.4 Memory Mask Field Encoding

**Memory and I/O Control Register for Bank 7..0
('MemControlReg(7..0)'),**

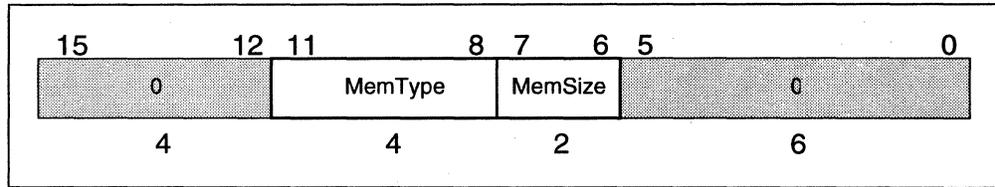


Figure 8.6 Memory and I/O Control Register Bit Assignments.

Bit	Assignment
11:8	Memory Type ('MemType')
7:6	Port Size Width ('MemSize')

Table 8.5 Memory and I/O Control Register Bit Assignments.

Memory Type ('Type') Field

The Type field determines the type of timing the Bus Interface will use. Values and actions for this field are listed in Table 8.6.

Value	Action
'1011'	PCMCIA-Memory-Style
'1010'	PCMCIA-I/O-Style
'1001'	M-Type I/O
'1000'	I-Type I/O
'0010'	Memory-Type for FCT260 (default)
'0001'	Memory-Type for FCT245
'0000'	Memory-Type for FCT543
All others	Reserved, undefined
<p>Note:</p> <ol style="list-style-type: none"> 1. PCI-Style supports a PCI subset that is likely to be used with PCI Peripheral chips. Both PCI master and slave (via external DMA) modes are supported. However, full support if needed must be externally provided, for instance, parity and terminate/stop issues. 2. PCMCIA-Style supports a PCMCIA host mode subset that is likely to be used with PCMCIA peripherals. PCMCIA-Memory and -IO Styles are intended for dynamic swapping by the software onto the same pair of chip selects. Typically, the Memory-Style is left on, and the I/O-Style is swapped in whenever it is needed, then swapped back to Memory-Style. 	

Table 8.6 Memory Type Field ('MemType') Encoding.

Port Size Width ('MemSize') Field

The PortSize field determines the width of the memory or I/O port. The value is inverted relative to the reset initialization vector value. Encoding information for this field is listed in Table 8.7

Value	Action
'11'	64-bit (32-bit 2-way interleaved) accesses (Valid for Memory Type only)
'10'	16-bit accesses
'01'	8-bit accesses
'00'	32-bit accesses

Table 8.7 PortSize ('MemSize') Encoding.

**Memory LSB Wait State Register for Bank 7..0
(‘MemLSBWaitStateReg(7..0)’)**

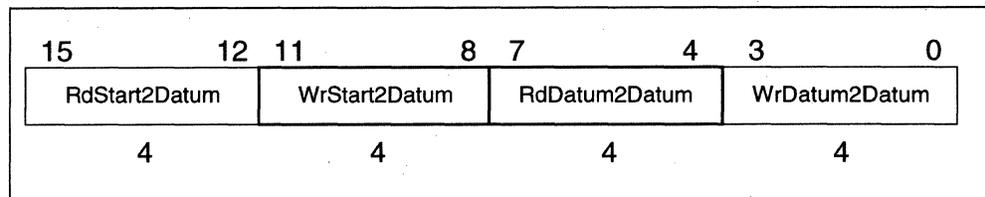


Figure 8.7 Memory LSB Wait State Register (‘MemLSBWaitStateReg’).

The Wait-State Generator registers provide fields to control the access timings to/from the CPU and the memory control and I/O control areas or these control areas for DMA accesses. The parameters controlled are:

1. Time from CS asserted to the first RdCEnN for read burst or time to RdCEnN and AckN for single word access or time to the first AckN for a burst DMA write.
2. Time between the RdCEnN's for burst reads or AckN's for burst DMA writes.
3. Time from the first RdCEnN to the AckN for burst reads or time from first AckN to last AckN for burst DMA writes.
4. The functionality of the SysWaitN signal for the corresponding control area.

There is a MemLSBWaitStateReg for each of the control areas or a total of 8 MemLSBWaitStateReg's.

The various fields and bit assignments of the Memory LSB Wait State Register are shown in Figure 8.7 and Table 8.8.

Bit	Assignment
15:12	RdStart2Datum
11:8	WrStart2Datum
7:4	RdDatum2Datum
3:0	WrDatum2Datum

Table 8.8 Memory LSB Wait State Register (‘MemLSBWaitStateReg’) Bit Assignments.

Read Start Cycle to the First Datum ('RdStart2Datum') Field:

and

Write Start Cycle to the First Datum ('WrStart2Datum') Field:

This field sets the number of cycles from the last (if repeated) 'S0 Start' bus cycle to the first RdCEnN for a burst read or to the RdCEnN and AckN for a single word read or to the first AckN for a burst DMA write or to the AckN for a single word write. The time can be from 1 cycle (Start2Datum = b'0000) (default) to 16 cycles (Start2Datum = b'1111). The wait states (Start2Rd value - 1) are injected onto the "2nd" clock edge after the bus cycle begins, such that the Second, 'S1' state is repeated. Encoding information for this field is contained in Table 8.9.

Value	Action
'15'	15 clock cycles to 1st datum (default).
...	
'0'	1 clock cycle from Start to 1st datum.
Note: At least 1 clock cycle is always implied by 'S1' state.	

Table 8.9 Start to the first Datum ('RdStart2Datum' and 'WrStart2Datum') Field Encoding.

Read Datum to Datum ('RdDatum2Datum') Field:

and

Write Datum to Datum ('WrDatum2Datum') Field:

This field sets the number of cycles between RdCEnN Datum for burst reads or AckN Datum for burst DMA writes. The time can be from 1 (Datum2Datum = b'0000) to 16 cycles (Datum2Datum = b'1111) such that the 'S2' state is repeated. Interleaved delay is between pairs of datum. Encoding information for this field is listed in Table 8.10.

Value	Action
'2' to '15'	3 to 16 cycles to next burst datum.
'1'	2 cycles to next burst datum (default)
'0'	1 cycle to next burst datum.

Table 8.10 Datum-to-Datum (RdDatum2Datum, WrDatum2Datum) Field Encoding

Memory MSB Wait State Register for Bank 7..0 (**MemMSBWaitStateReg(7..0)**)

These registers are both readable and writable. Programming information for this field is located in Figure 8.8 (Memory MSB Wait State Register) and Table 8.11 (Memory MSB Wait State Register Bit Assignments).

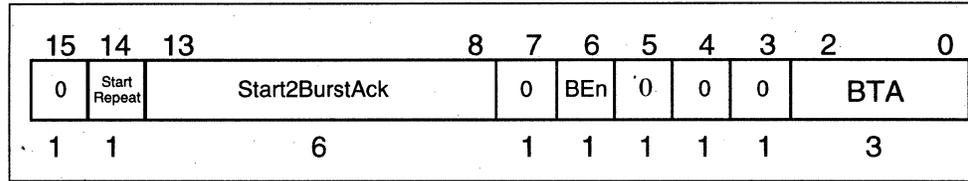


Figure 8.8 Memory MSB Wait State Register ('MemMSBWaitStateReg').

Bit	Assignment
15	Reserved to '0'
14	StartRepeat (default is 1)
13:8	Start2BurstAck (default is 0)
7	Reserved to '0'
6	BEn (default is 0)
5:3	Reserved to '0'
2:0	BTA (default is 0)

Table 8.11 Memory MSB Wait State Register ('MemMSBWaitStateReg') Bit Assignments.

Repeat Start Bus Cycle State 0 ('StartRepeat') Field

This field controls the number of times the S0 first bus cycle state is repeated until the second 'S1' state is entered. An application is to allow more time for address setup to the chip select on the slowest 600ns PCMCIA cards. Field encoding information for this field is located in Table 8.12.

Value	Action
'1'	Repeat Start Cycle 1 time (default).
'0'	Repeat Start Cycle 0 times (typical).

Table 8.12 Repeat Start Bus Cycle State 0 ('StartRepeat') Field Encoding.

Start of Read to AckN on Burst Reads ('Start2BurstAck') Field

This field sets the number of cycles from the first bus cycle to the AckN for a burst read or to the last AckN for a burst DMA write. The time can be from 0 cycles (Start2BurstAck = b'000000) to 62 cycles (Start2BurstAck = b'111110). This field is only valid for Memory type and not valid for I/O type accesses. Field encoding information for this field is listed in Table 8.13.

Value	Action
'63'	No Ack (default).
'62'...'0'	62 to 0 cycles until Ack (typical is 0).

Table 8.13 First Read to AckN on Burst Reads ('Start2BurstAck') Field Encoding.

Byte Enables on Reads ('BEn') Field

This field selects whether the $\overline{\text{MemWrEn}}(3:0)$ pins are asserted on reads. It is recommended that byte enables only be enabled for reads on PCI accesses in order to support future Fly-By DMA modes on other peripherals that may use $\overline{\text{MemWrEn}}(3:0)$. $\overline{\text{MemWrEn}}(3:0)$ always return high at the end of a bus transaction. Field encoding information for this field is listed in Table 8.14.

Value	Action
'1'	Allow Byte Enables on reads.
'0'	Allow Byte Enables on writes only with $\overline{\text{MemWrEn}}$ (default).

Table 8.14 Byte Enables on Reads ('BEn') Field Encoding.

Bus Turn-Around ('BTA') Field

This field sets the minimum number of idle data bus cycles after a read. This field must be set to at least '1' because of the possibility of a read followed by a write. On a read followed by a read, the number of idle cycles is from the pertinent read enable de-asserting until the next pertinent read enable (Memory, I/O, or DRAM) asserting. Field encoding information for this field is listed in Table 8.15

Value	Action
'111'	Minimum of 7 clocks (may increase in future products)
'110'	6 clocks
'101'	5 clocks
'100'	4 clocks
'011'	3 clocks
'010'	2 clocks
'001'	1 clock (default)
'000'	Reserved

Table 8.15 Bus Turn-Around ('BTA') Field Encoding.

Memory Controller Timing Diagrams

This section illustrates a number of timing diagrams applicable to R36100 RISController memory transactions. The values for the AC parameters are contained in the separate document, "R36100 RISController Data Sheet."

Read Transactions

The bus interface timing for read transactions is described in the following section. The internal bus interface to CPU core timing for reads is described in Chapter 7, "Bus Interface Unit Controller".

Basic 1-Datum Read with 0 Wait-States

Figure 8.9 illustrates a basic Memory-Type Memory Controller read transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysRd}}$ and the appropriate $\overline{\text{MemCSj}}$ assert (if they are not already in this state, as the result of a previous transaction) to indicate the read transaction and which memory bank is being used. After the completion of this initial 'Start' cycle stage, the Data Sampling stage is begun. In this second stage, the appropriate Memory Read Enable signal, either $\overline{\text{MemRdEnEven}}$ or $\overline{\text{MemRdEnOdd}}$, will assert to allow the external memory bank to turn on and begin driving data back to the RISController.

Since Figure 8.10 is for a single datum read, this Data Sampling Stage is the last Datum for this transaction and thus $\overline{\text{SysBurstFrame}}$ is de-asserted. To end a Data Sampling Stage, $\overline{\text{SysDataRdy}}$ asserts on clock cycles where data is expected. After the last Data is sampled, note that the signals, $\overline{\text{SysRd}}$ and $\overline{\text{MemCS}}$ may not necessarily de-assert as the next bus transaction may already be starting. The internal signal, AckN , that is associated with R3051-family read cycles is generated automatically for single word reads (and writes).

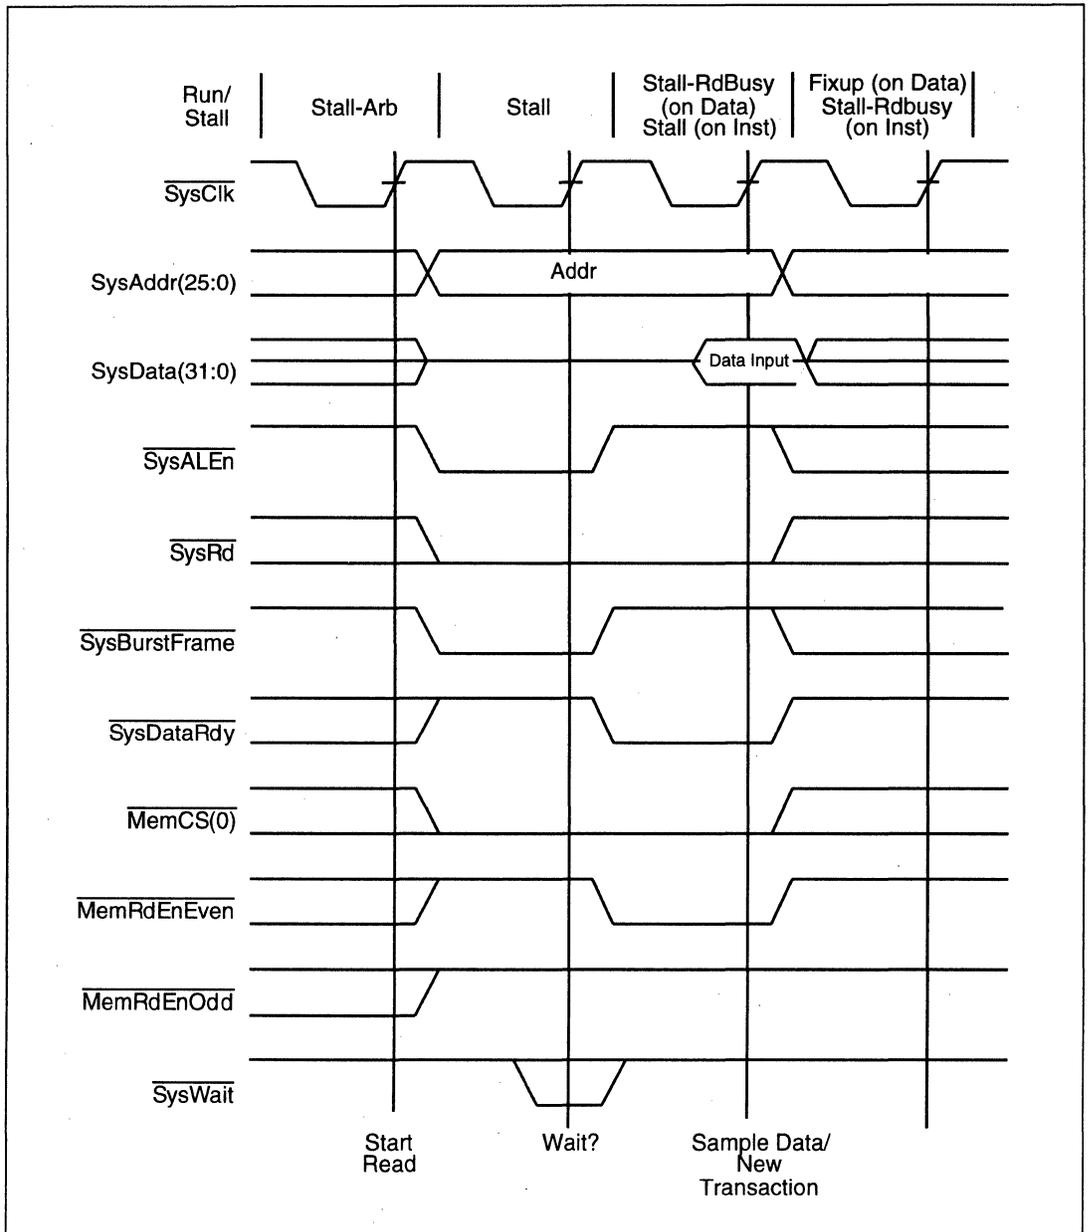


Figure 8.9 1-Datum Read with 0 Wait-States.

1-Datum Read with 0 Wait-States using Odd Chip Select

Figure 8.10 illustrates a basic Memory-Type Memory Controller read transaction, except that the access is to an Odd Memory Chip Select, MemCS(7,5,3,1) instead of an even one. Because an odd MemCS is asserted, MemRdEnOdd is asserted during the Data Sampling Stage instead of its even counterpart.

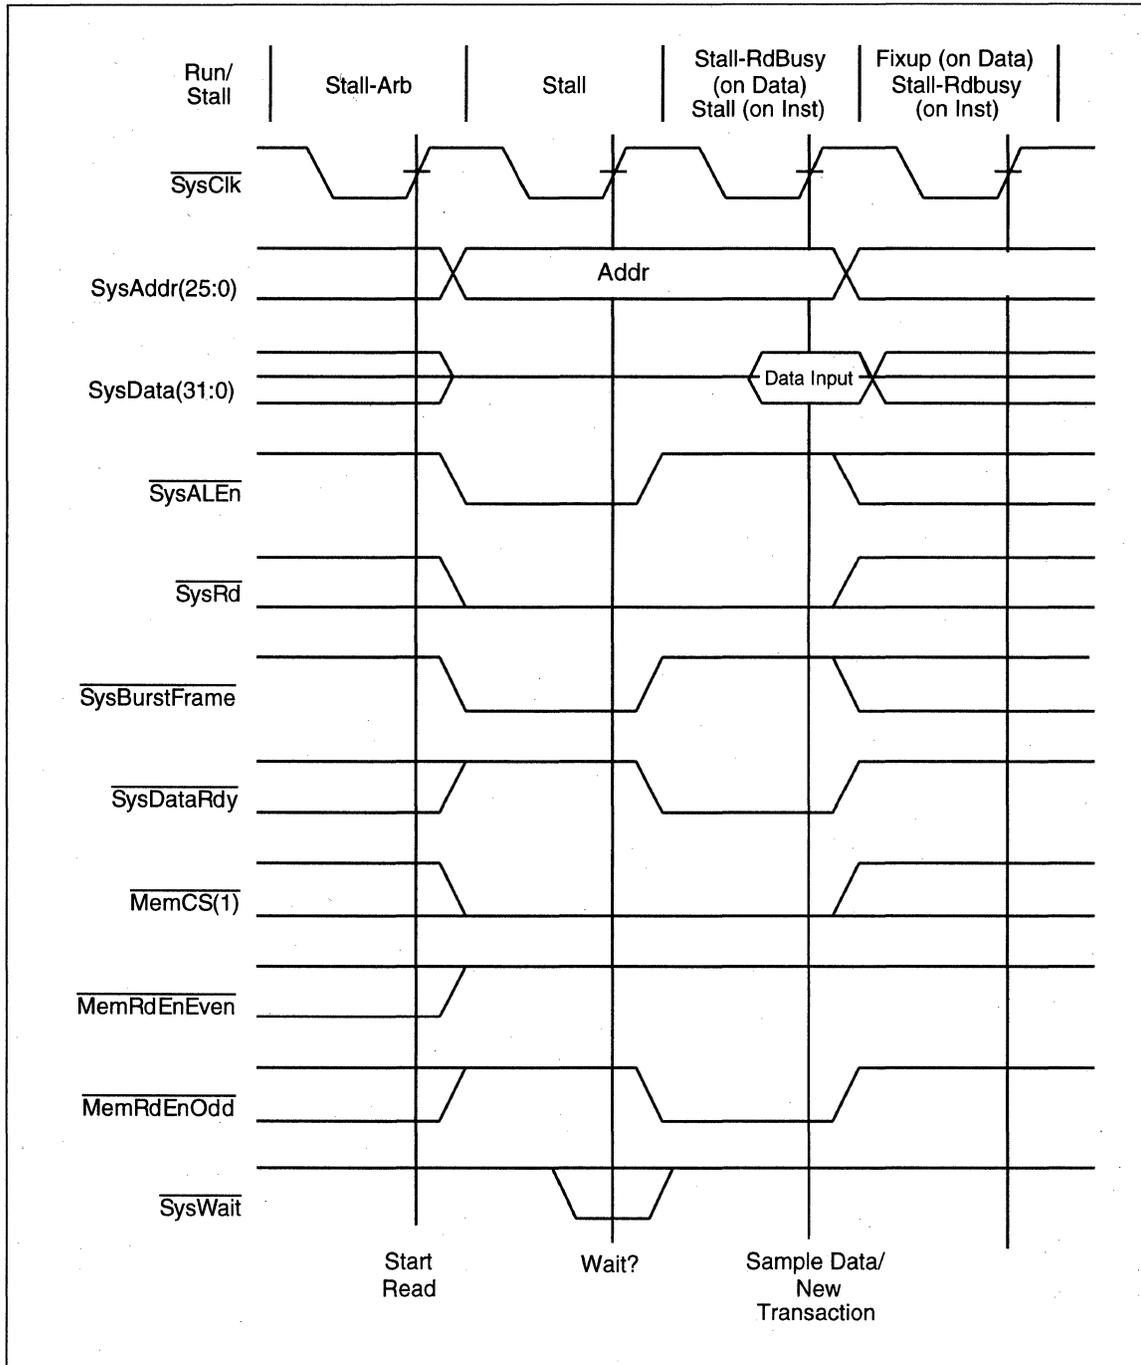


Figure 8.10 1-Datum Read with 0 Wait-States Using an Odd Chip Select.

Read with Wait-State using Start Repeat Field

Figure 8.11 illustrates a basic Memory-Type Memory Controller read where 1 wait-state has been added by repeating the Start Cycle. This special effect is programmed into the Wait-State Generator using the Start Repeat Field in the MemLSBWaitStateReg() Register. When the Start Cycle repeats, the Data Sampling Stage is delayed and the assertion of the Memory Read Enable strobe, MemRdEnEven or MemRdEnOdd is delayed. This effect is useful for very slow memories or memories that require significant address setup before the chip is selected. An example is the 600ns access time mode of the PCMCIA memory protocol. The Start repeat Field affects both reads and writes.

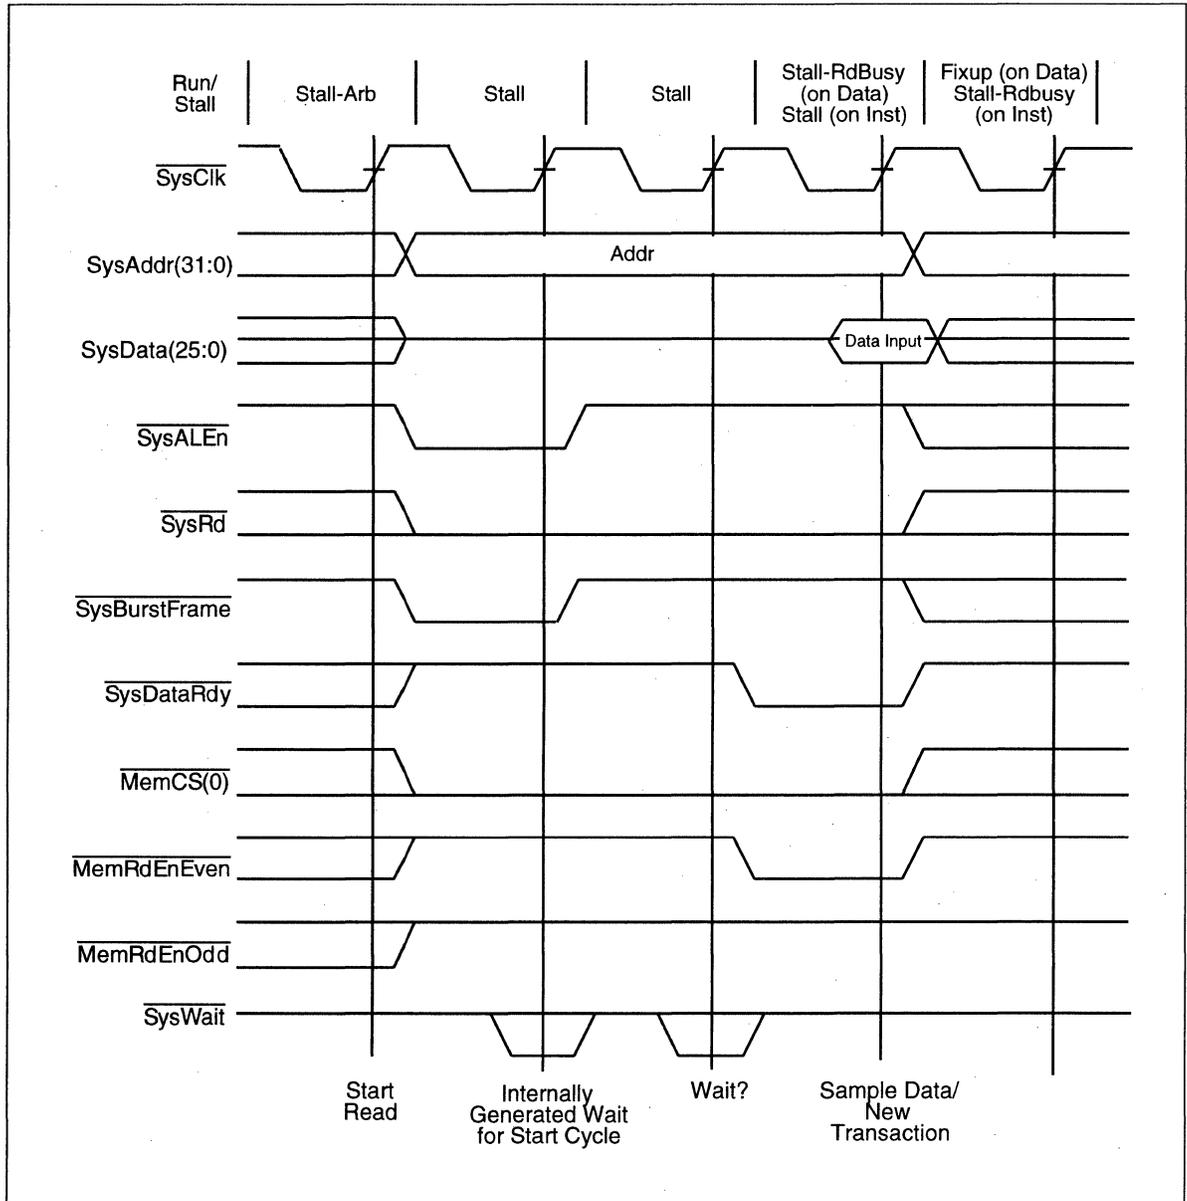


Figure 8.11 1-Datum Read with 1 Wait-State using StartRepeat Field.

Read with Wait-State using RdStart2Datum Field

Figure 8.12 illustrates a basic Memory-Type Memory Controller read where 1 wait-state is added using the RdStart2Datum Field of the MemMSBWaitStateReg0 Register. Any number from 0 to 15 internal wait-states may be added using the RdStart2Datum Field. With this field, the Memory Read Enable strobe, either MemRdEnEven or MemRdEnOdd is asserted as normal, but then wait-states are added where SysDataRdy is not asserted until the RdStart2Datum Field has finished counting. When SysDataRdy is asserted, then the Data from the external Memory Bank is sampled into the RISController.

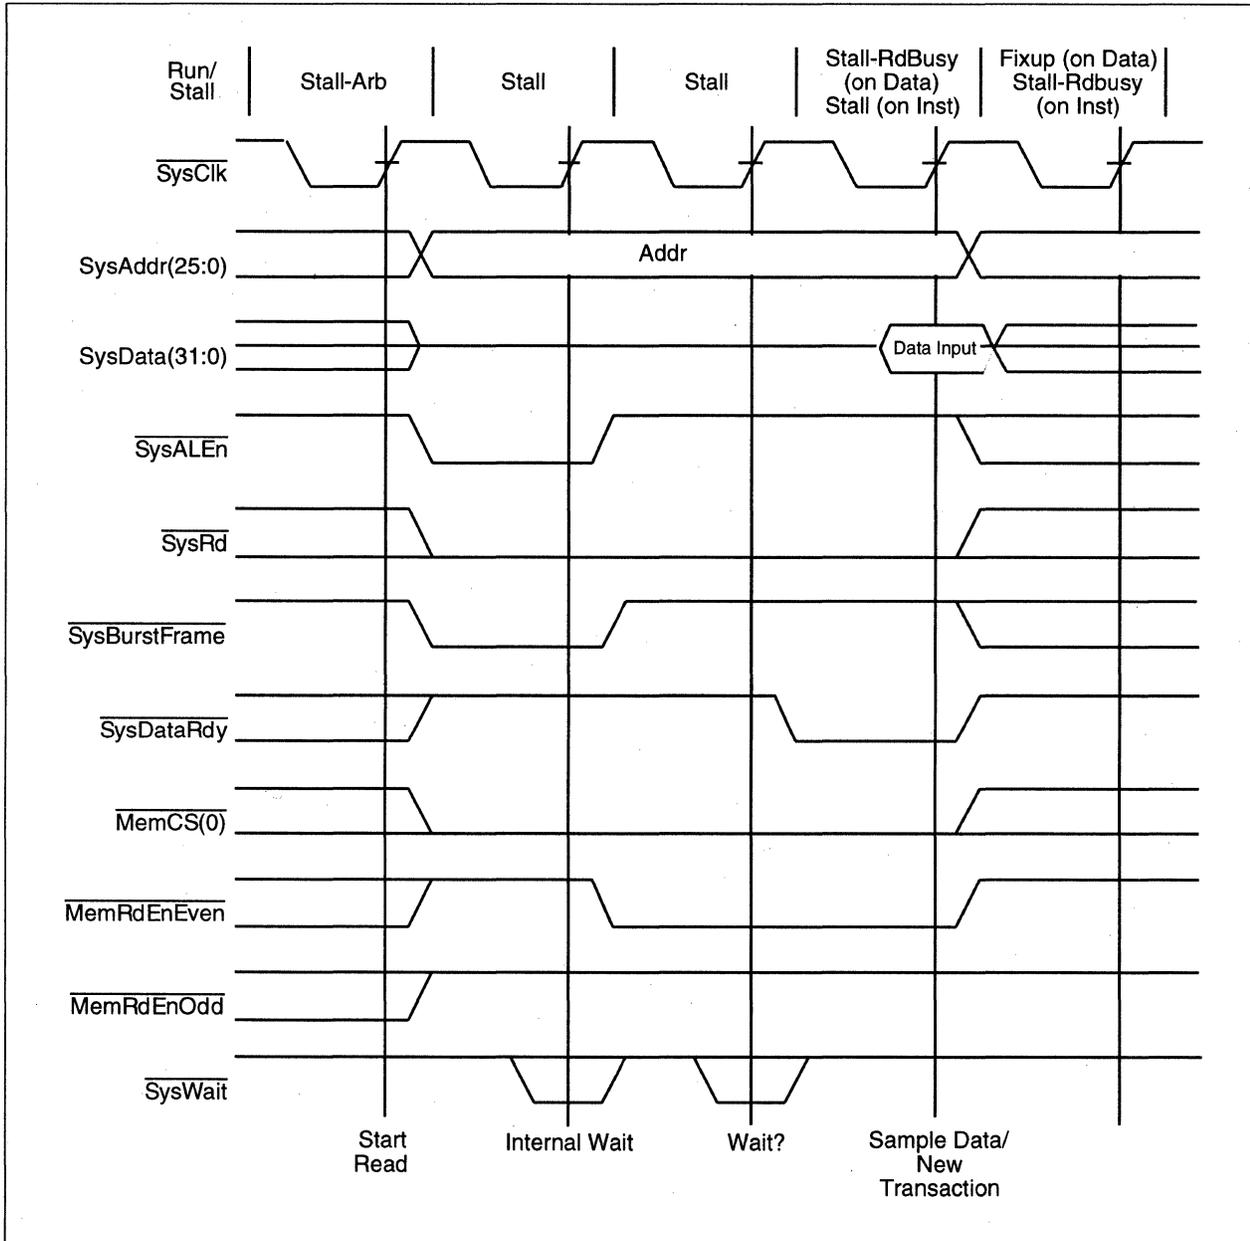


Figure 8.12 Read with Wait-State using RdStart2Datum Field.

Read with Wait-State using SysWait

Figure 8.13 illustrates a basic Memory-Type Memory Controller read where 1 wait-state is added using the external signal pin, SysWait. SysWait is not expected to be used for conventional memories, since it is easier to program the Wait-State Generator to produce wait-states. However, SysWait can be useful for Dual-Port memory and off-card memories where there may be an indeterminate amount of time before the access can begin. Since SysWait is sampled a clock ahead of when it is used, its effect is seen two clocks later than when it is asserted. If SysWait is asserted when SysDataRdy is asserted then an additional Data Sampling clock cycle is repeated with SysDataRdy remaining low. Thus external logic analyzers or other debug equipment may want to gate SysDataRdy with SysWait in order to decode valid Data samples.

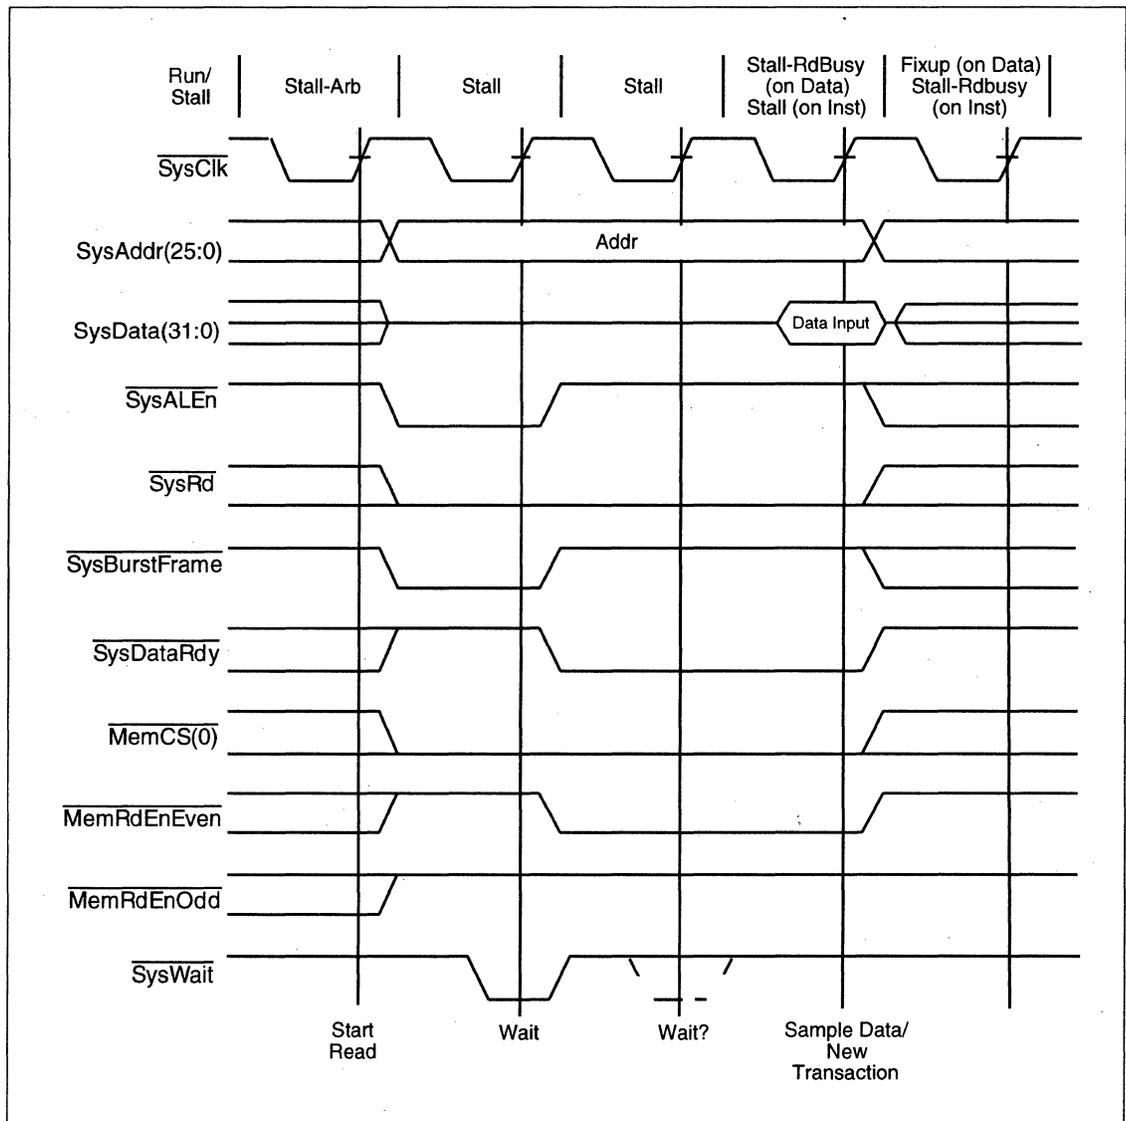


Figure 8.13 Read with Wait-State using SysWait.

4-Word Burst Read with 0 Wait-States

Figure 8.14 illustrates a 4-word Burst Read when using the Memory-Type Memory Controller. This example can also be generalized to demonstrate any multi-datum read generated from the CPU core which includes double, triple, quad, and 16-byte reads using an 8-bit port width, and double, and octi-halfword reads using a 16-bit port width. Note that the DMA Controller can potentially do any length from 1 to 16 datum.

As shown in Figure 8.14, SysBurstFrame can be used to determine which is the last datum to be sampled. After the first datum is sampled, the Wait-State Generator uses the RdDatum2Datum field in the MemMSBWaitStateReg() Register in order to determine the number of wait-states to generate. Figure 8.15 gives an example of inserting wait-states between later data elements. By programming different values into the RdStart2Datum and the RdDatum2Datum Fields, a burst read can be throttled to, for example, give a longer access time to the first datum and shorter access time for any subsequent datum.

To obtain the maximum optimization of a 4-word burst read (4 32-bit datum, 8 16-bit datum or 16 8-bit datum) the Start2BurstAck Field of the MemMSBWaitStateReg() Register must be programmed. Programming this field to 3 clock cycles before the last datum is sampled places the internal R3051-family like signal, AckN, such that the CPU pipeline can be restarted early. Systems that have an indeterminate number of external SysWait wait-states must program this field to give an internal AckN. In such cases, like single word reads, AckN is generated automatically on the last datum sample.

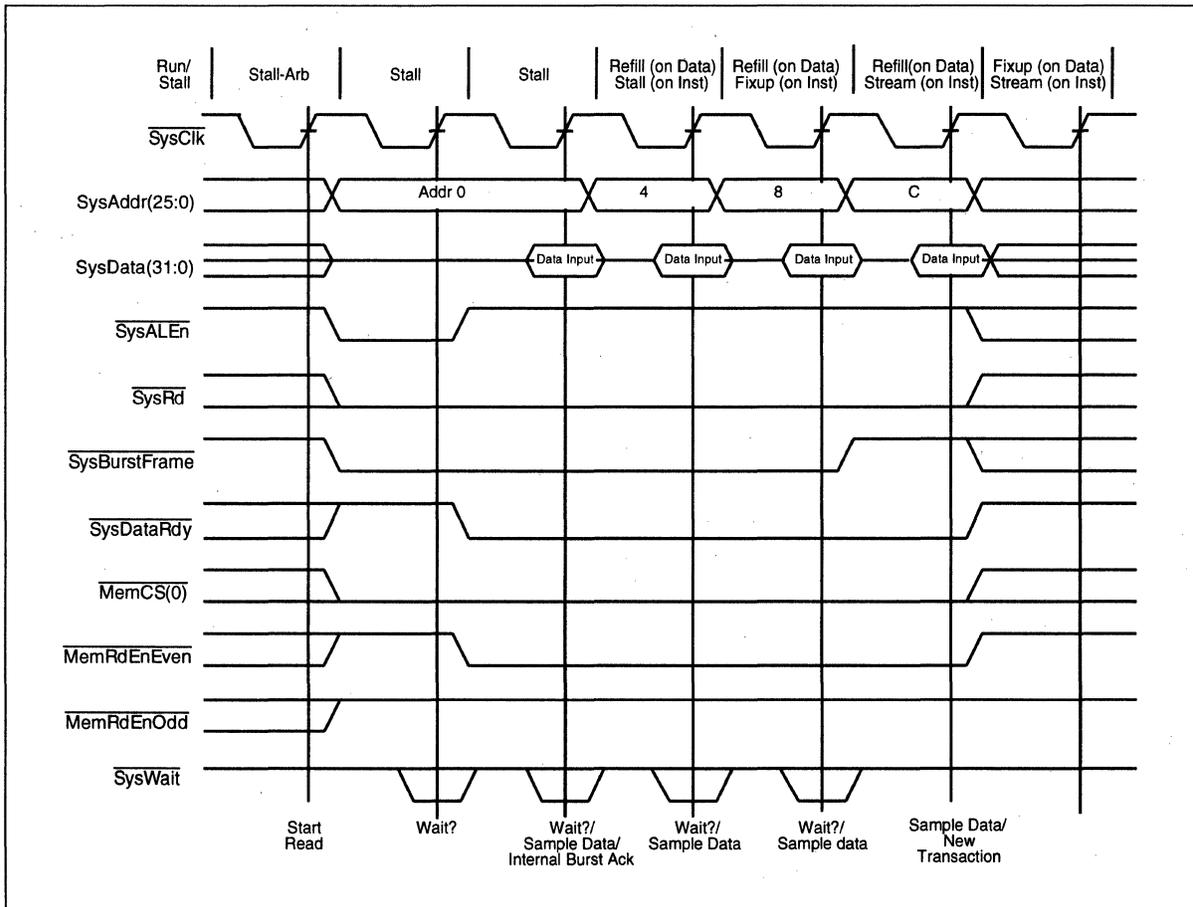


Figure 8.14 4-Word Burst Read with 0 Wait-States.

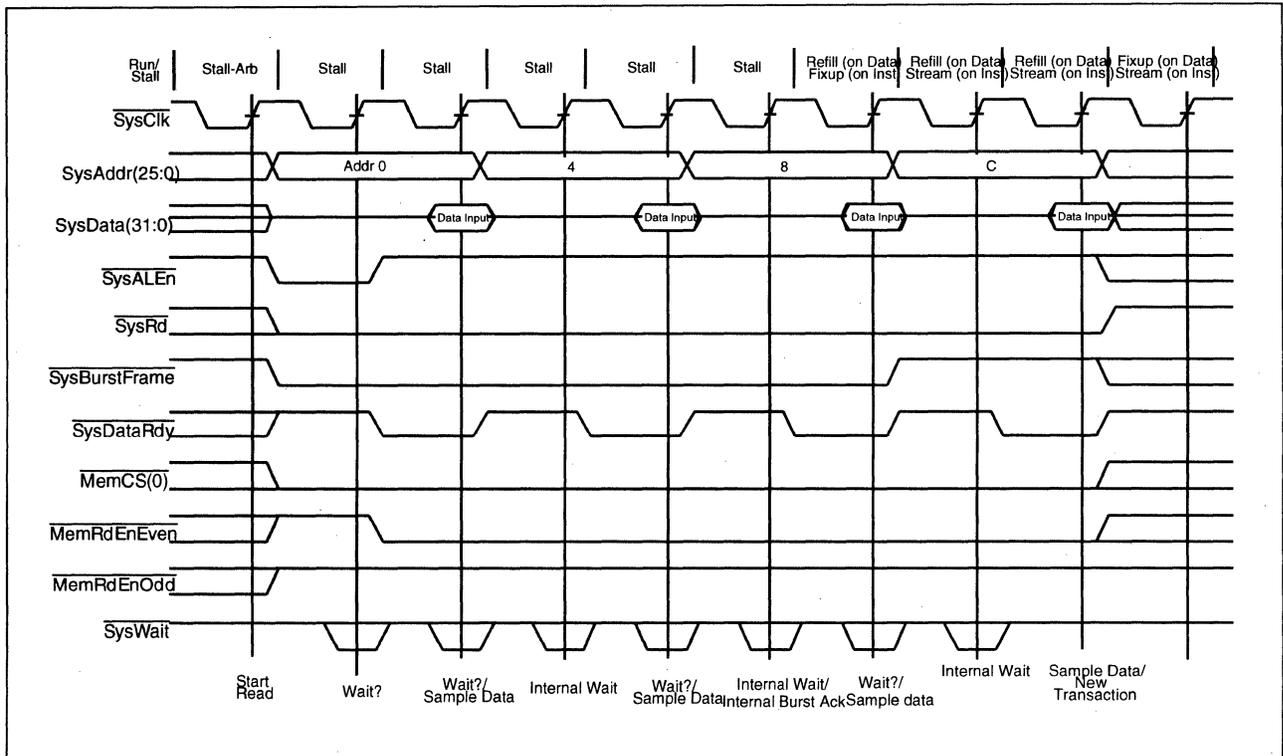


Figure 8.15 4-Word Burst Read with Wait-States using RdDatum2Datum Field.

Basic 16-bit PCMCIA-style I/O Read with 0 Wait-States

Figure 8.16 illustrates a basic 16-bit PCMCIA-Style Memory Controller read transaction. Each transaction begins with both SysALEn and SysBurstFrame asserting. At this time, SysRd asserts (if it is not in this state, as the result of the previous transaction). Assuming there are no internally programmed StartRepeat wait-states, on the next clock cycle, SysBurstFrame de-asserts and the MemCS(0) pair asserts. Note that on PCMCIA transactions the MemCS(0) pair is asserted, according to which bytes are enabled and valid. Thus if the even byte is valid, then the even MemCS(0) will assert. If the odd byte is valid, then the odd MemCS(0) will assert. If both bytes are valid, then both MemCS(0)s in the pair assert. Assuming there are no internally programmed RdStart2Datum wait-states, SysDataRdy asserts to indicate that the data from the memory device is being sampled into the RISController. On the next clock--the final clock of the transaction--the MemCS(0) pair de-asserts and, simultaneously, the next transaction may begin.

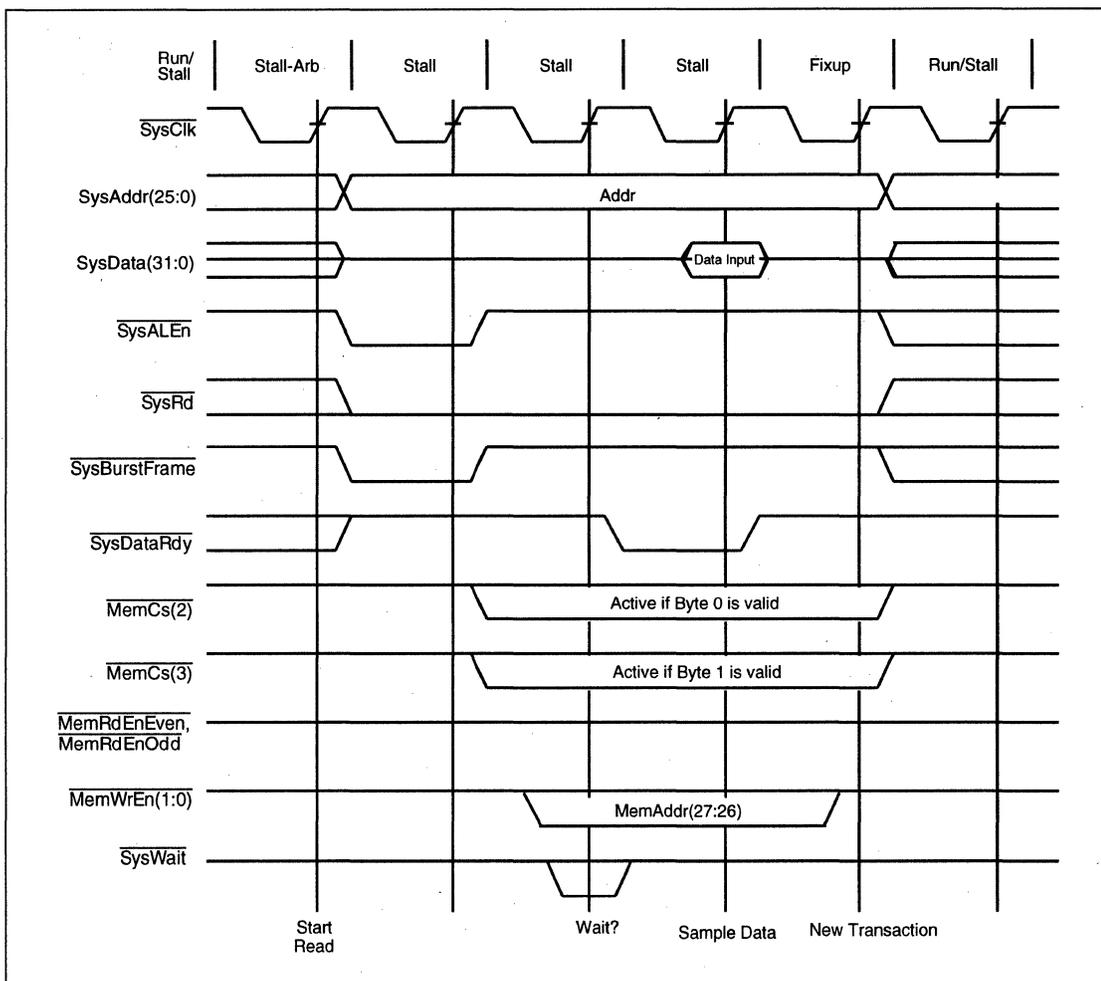


Figure 8.16 PCMCIA-Style Memory Read with 0 Wait-States

Write Transactions

The bus interface aspect of write transactions is described in the following section. The internal bus interface to CPU core aspect of writes is described in Chapter 7, “Bus Interface Unit Controller”.

Single Datum Write

Figure 8.17 illustrates a basic Memory-Type Memory Controller write transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysWr}}$ and the appropriate $\overline{\text{MemCS}}$ assert (if they are not already in this state) to indicate the write transaction and the memory bank being used. After completing this Start cycle stage, the Data Driving stage begins. In this second stage, $\overline{\text{MemWrEn}}(3:0)$ (which acts as write byte enable strobes) assert. In general, from 1 to 4 of the $\overline{\text{MemWrEn}}(3:0)$ signals will be asserted.

Since Figure 8.17 is for a single datum write, this Data Driving Stage is the last Datum for this transaction and thus $\overline{\text{SysBurstFrame}}$ is de-asserted. To end a Data Driving Stage, $\overline{\text{SysDataRdy}}$ asserts on clock cycles where data is expected to be latched with the trailing de-asserting edges of $\overline{\text{MemWrEn}}(3:0)$. After the last Data is driven, note that signals such as $\overline{\text{SysWr}}$ and $\overline{\text{MemCS}}$ may not necessarily de-assert, as the next bus transaction may already be starting. The internal signal, AckN , associated with R3051-family write cycles, is generated automatically for single word writes (and reads).

Note: This manual does not illustrate a Memory-type write transaction to an Odd Memory Chip Select. However, the only difference is that during the Data Driving Stage an odd $\overline{\text{MemCS}}$ is asserted instead of its even counterpart.

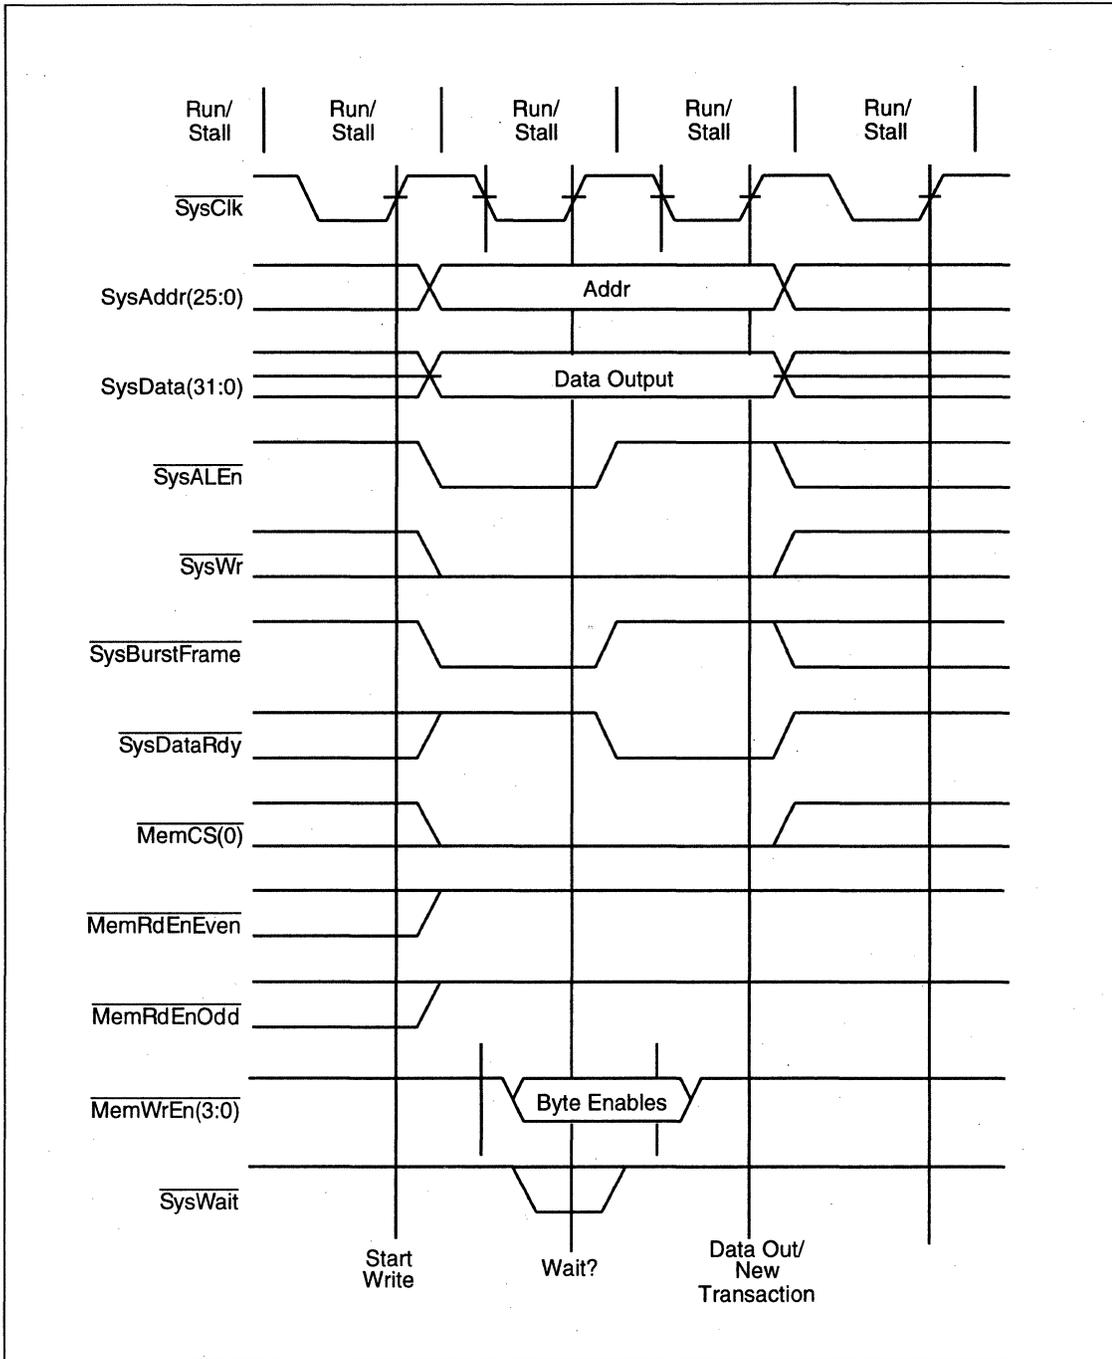


Figure 8.17 1-Datum Write with 0 Wait-States

1-Datum Write with 0 Wait-States using FCT245-Type Field

Figure 8.18 illustrates a basic write using the FCT245-Type Field. Either the even or the odd read enable, MemRdEnEven or MemRdEnOdd will assert on even or odd writes such that a FCT245 transceiver output enable can be connected.

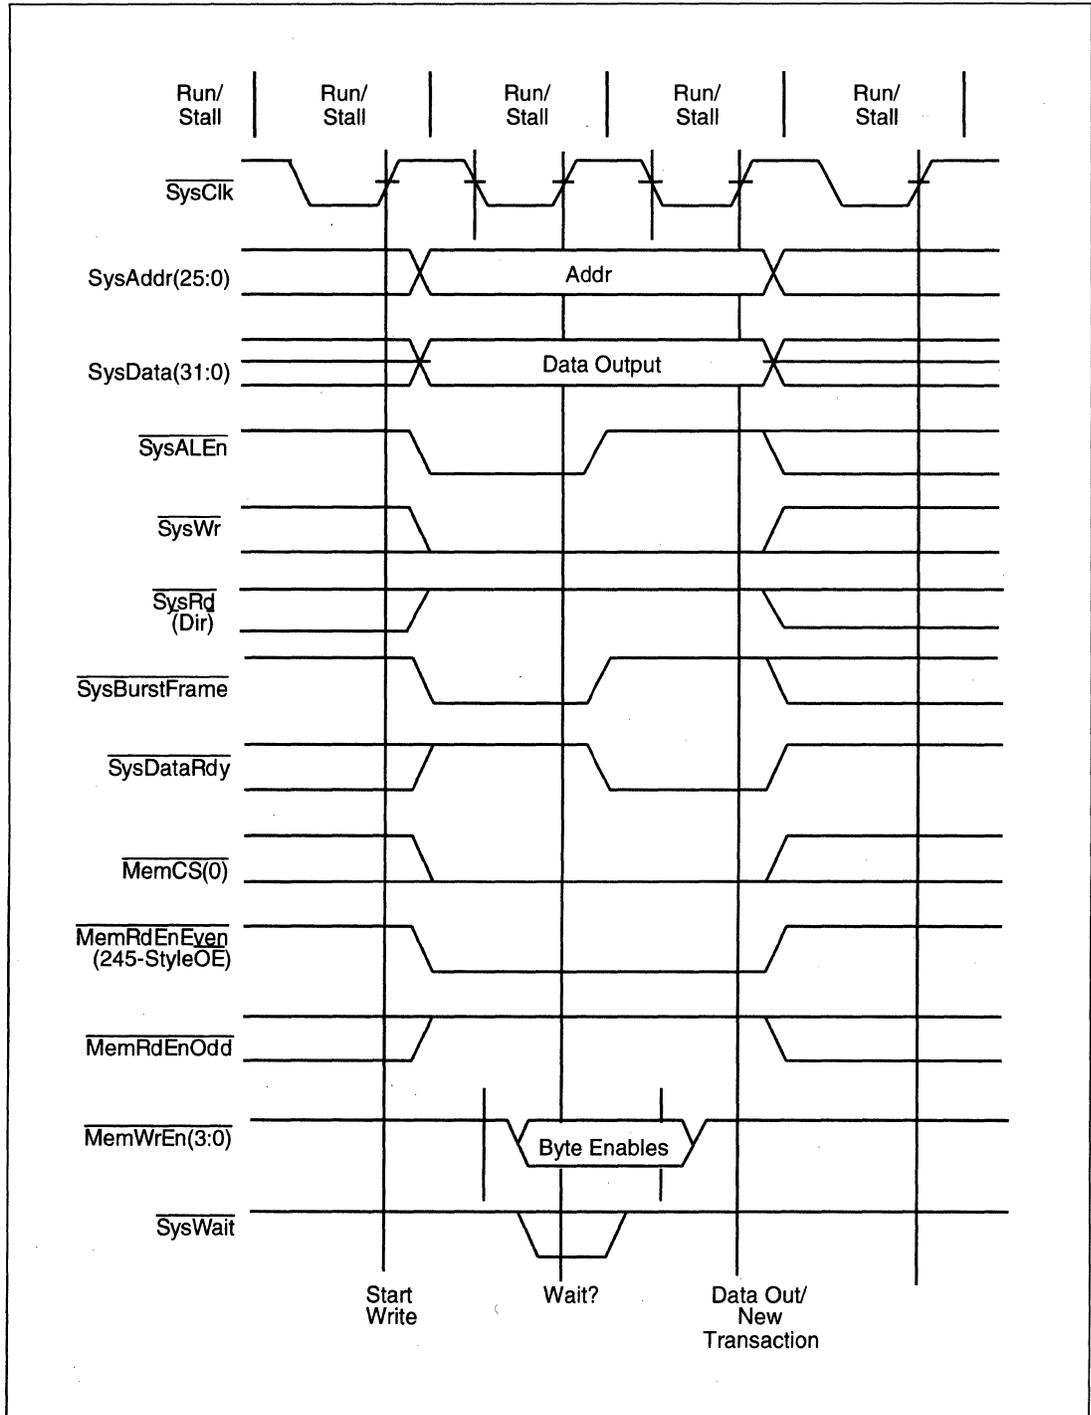


Figure 8.18 1-Datum Write with 0 Wait-States using FCT245-Type Field.

1-Datum Write with Wait-State using StartRepeat Field

Figure 8.19 illustrates a basic Memory-Type Memory Controller write where 1 wait-state has been added by repeating the Start Cycle. This special effect is programmed into the Wait-State Generator using the Start Repeat Field in the MemLSBWaitStateReg0 Register. When the Start Cycle repeats, the Data Driving Stage is delayed and the assertion of the Memory Write Enable strobes, MemWrEn(3:0), are delayed. This effect is useful for very slow memories or memories that require significant address setup before the chip is selected. For example, with the 600ns access time mode of the PCMCIA memory protocol, the Start repeat Field affects both reads and writes.

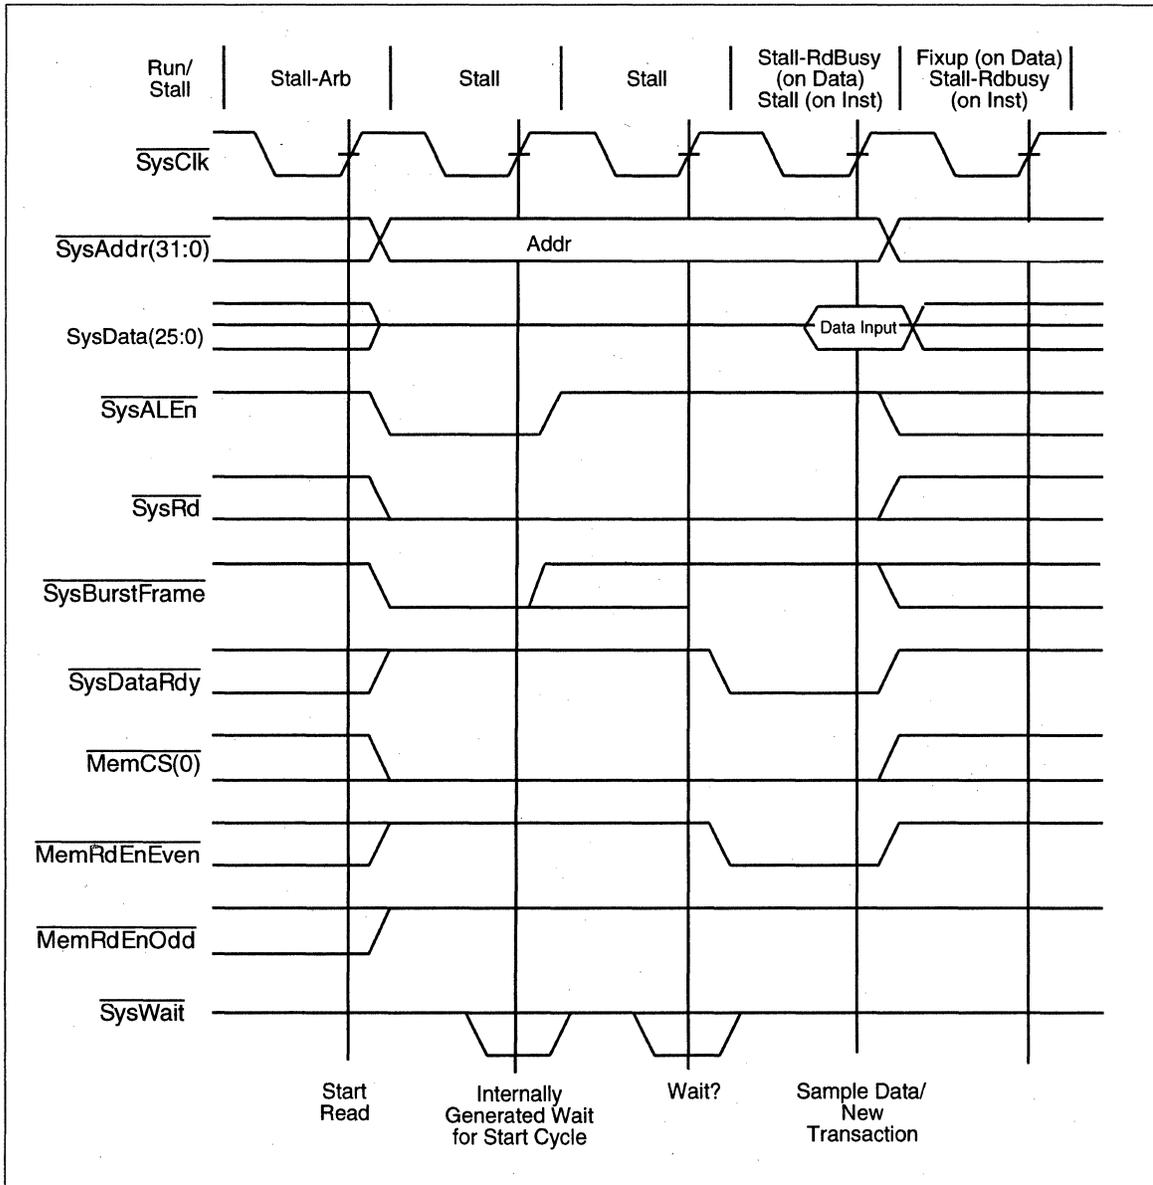


Figure 8.19 1-Datum Write with Wait-State using StartRepeat Field.

1-Datum Write with Wait-State using WrStart2Datum Field

Figure 8.20 illustrates a basic Memory-Type Memory Controller write where 1 wait-state is added using the WrStart2Datum Field of the MemMSBWaitStateReg() Register. Any number from 0 to 15 internal wait-states may be added using the WrStart2Datum Field. With this field, the Memory Write Enable strobe (either MemWrEnEven or MemWrEnOdd) as well as the write byte enables (MemWrEn(3:0)) are asserted as normal. However, wait-states are added (where SysDataRdy is not asserted) until the WrStart2Datum Field has finished counting. When SysDataRdy is asserted, then the Data from the external Memory Bank is sampled by external memory.

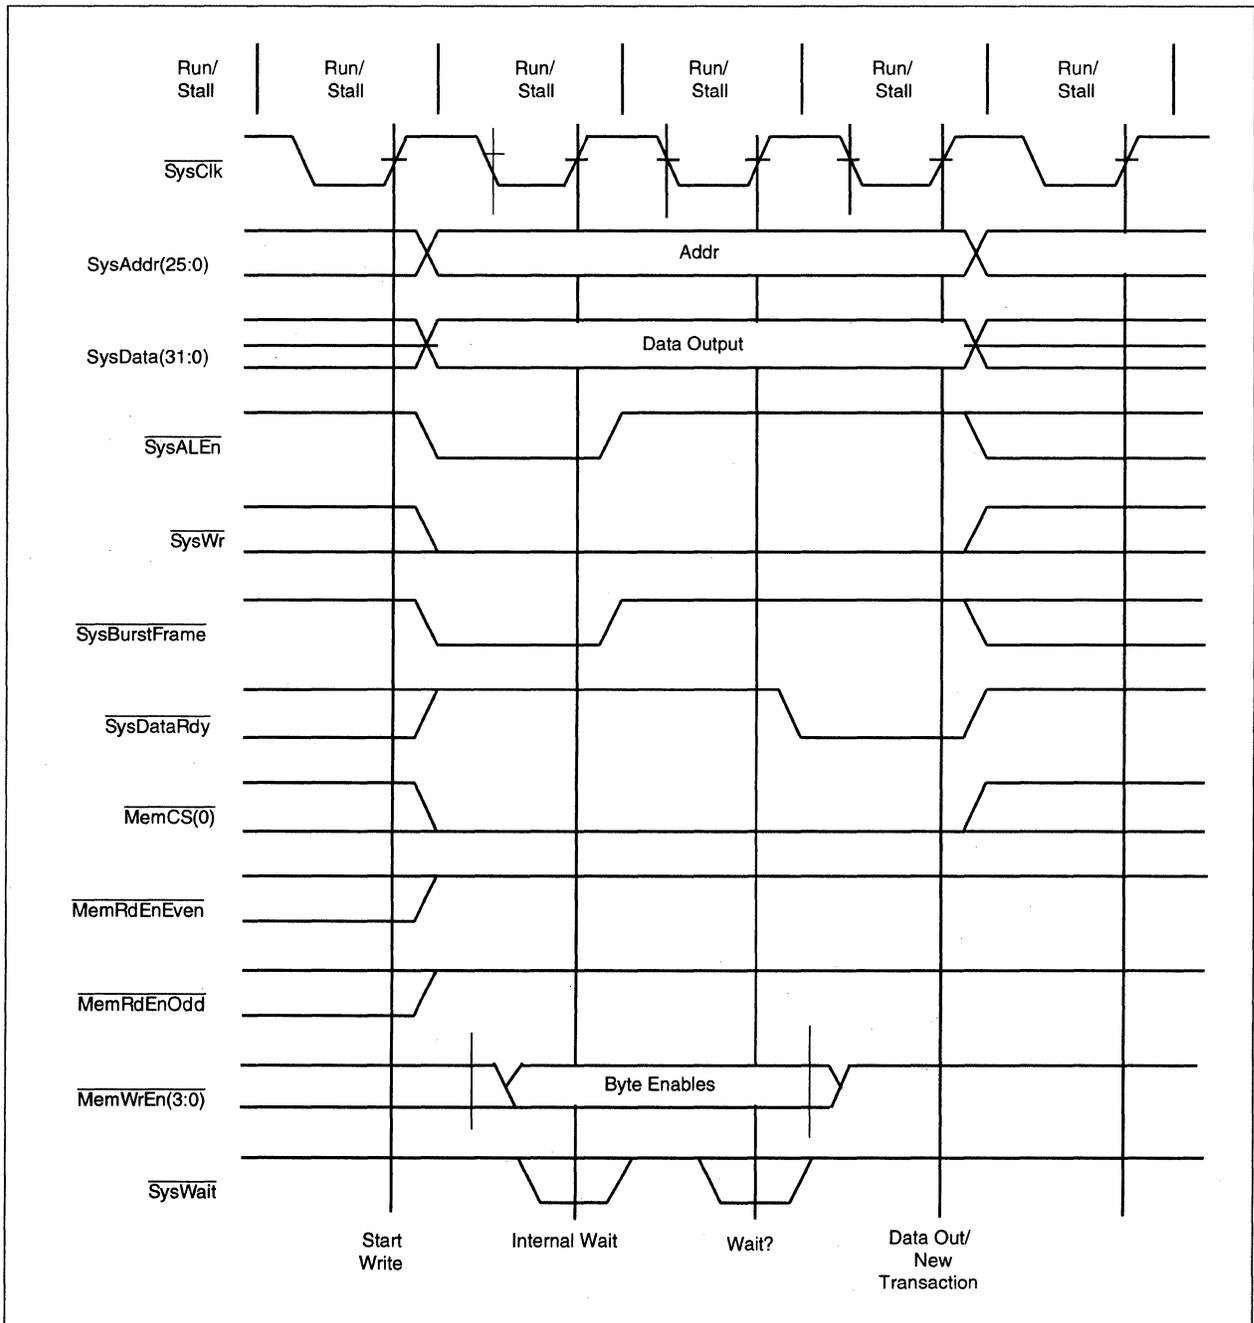


Figure 8.20 1-Datum Write with Wait-State using WrStart2Datum Field.

1-Datum Write with Wait-State using SysWait

Figure 8.21 illustrates a basic Memory-Type Memory Controller write where 1 wait-state is added using the external signal pin, SysWait. SysWait is not expected to be used for conventional memories, since it is easier to program the Wait-State Generator to produce internal wait-states. However, SysWait can be useful for Dual-Port memory and off-card memories where there may be an indeterminate amount of time before the access can begin. Since SysWait is sampled a clock ahead of when it is used, its effect is seen two clocks later than when it is asserted. If SysWait is asserted when SysDataRdy is asserted then an additional Data Sampling clock cycle is repeated with SysDataRdy remaining low. Thus external logic analyzers or other debug equipment may want to gate SysDataRdy with SysWait in order to decode valid Data samples.

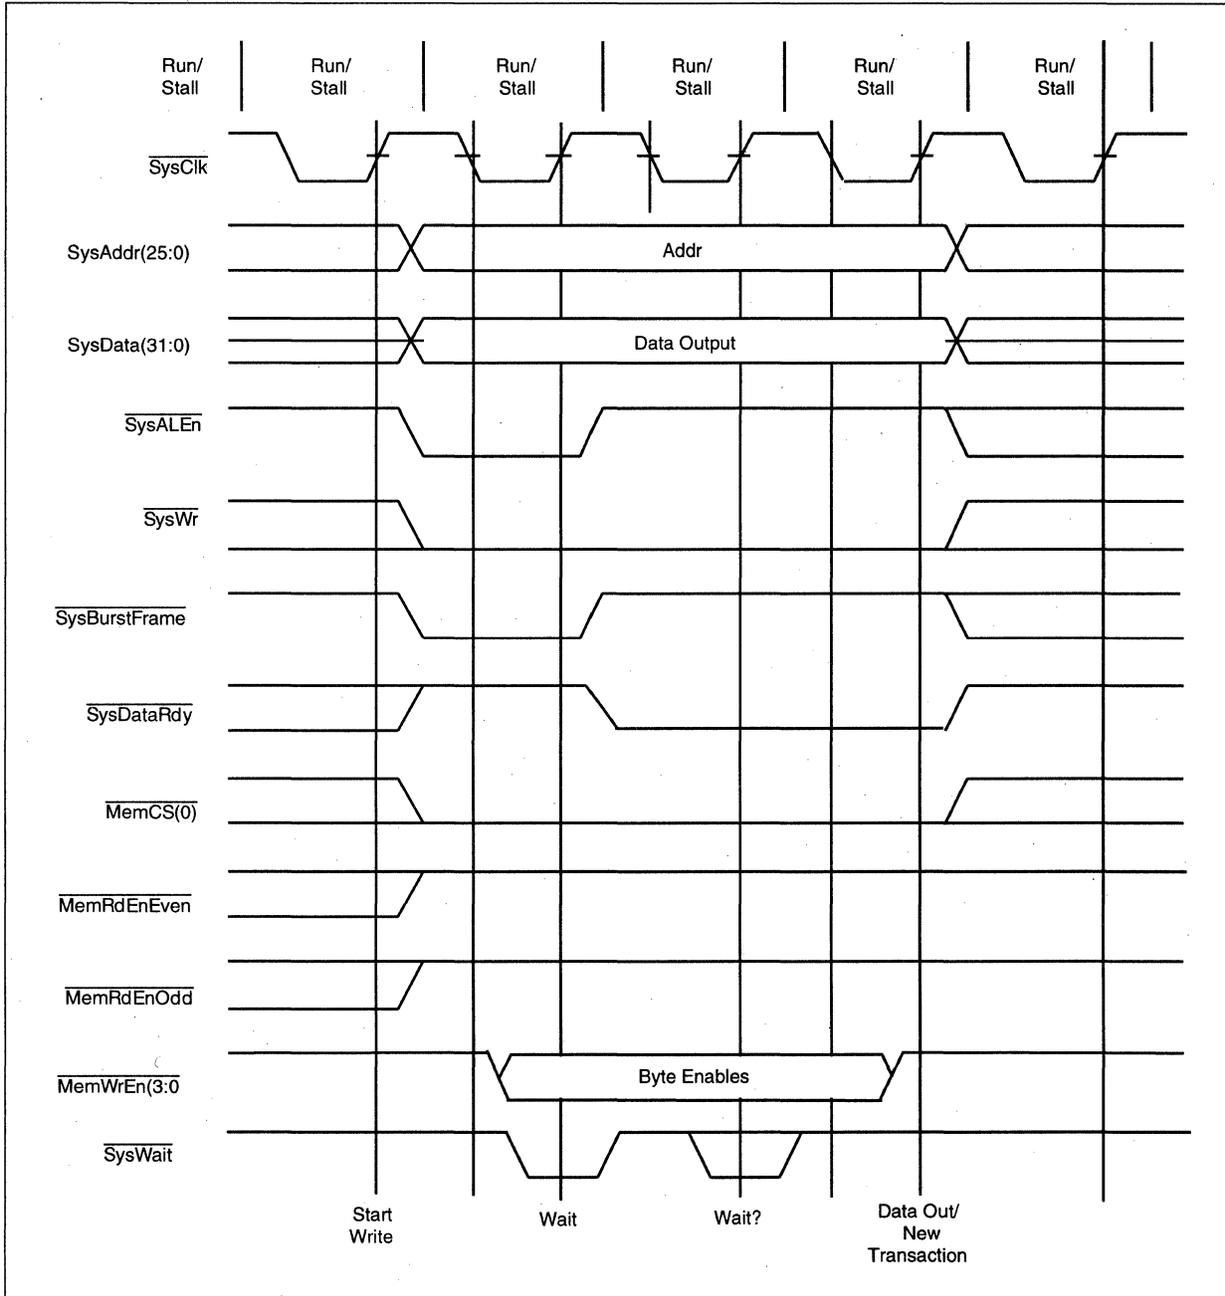


Figure 8.21 1-Datum Write with Wait-State using SysWait.

Multi-Datum Burst Write

Figure 8.22 illustrates a 2-datum Burst Write when using the Memory-Type Memory Controller. This example can also be generalized to demonstrate any multi-datum write generated from the CPU core which includes double, triple, and quad byte writes using an 8-bit port width, and double halfword writes using a 16-bit port width. Note that although the CPU core will not generate bursts in the 32-bit width, the internal and external DMA Controller channels can do so, and can potentially do any length.

In Figure 8.21, SysBurstFrame can be used to determine which is the last datum to be sampled. At the end of each datum, MemWrEn(3:0) de-asserts for 1 clock cycle. They re-assert if more datum needs to be processed. After the first datum is sampled, then the Wait-State Generator uses the WrDatum2Datum Field in the MemMSBWaitStateReg() Register to determine the number of wait-states to generate.

Most conventional SRAMs need to use WrDatum2Datum wait-states to de-assert MemWrEn(3:0) at the end of each word written. The zero wait-state case is useful for certain types of FIFOs.

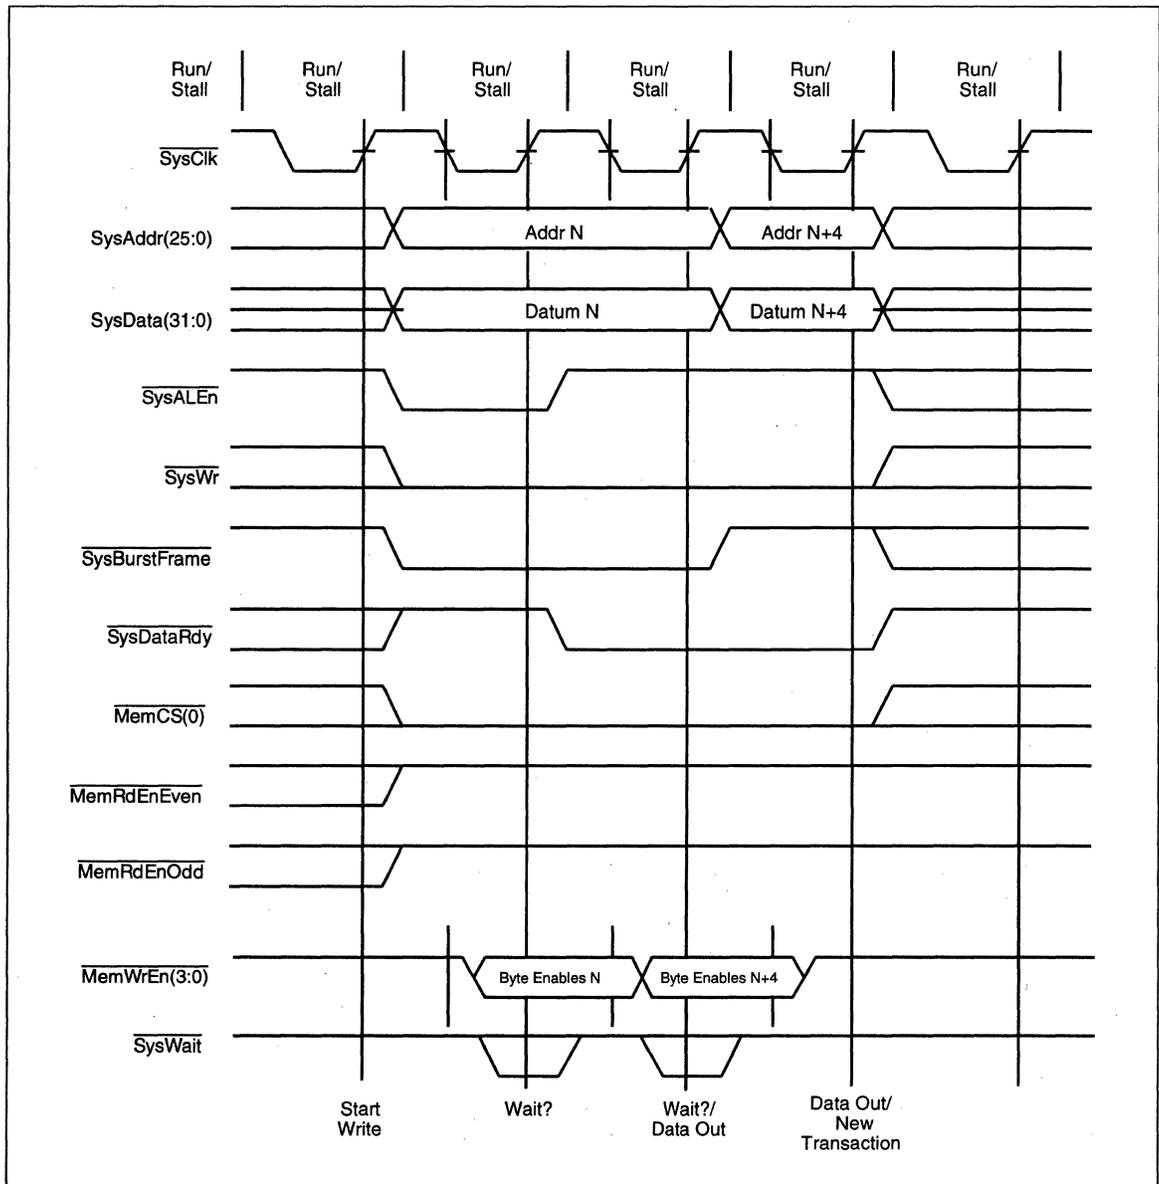


Figure 8.22 Multi-Datum Burst Write.

Multi-Datum Burst Write using Wait-State with WrDatum2Datum

Figure 8.23 gives an example of inserting wait-states between later data elements. By programming different values into the WrStart2Datum and the WrDatum2Datum Fields, a burst write can be throttled to for instance give a longer access time to the first datum and shorter access time for any subsequent datum.

Like single word writes, AckN is generated automatically on the last datum sample.

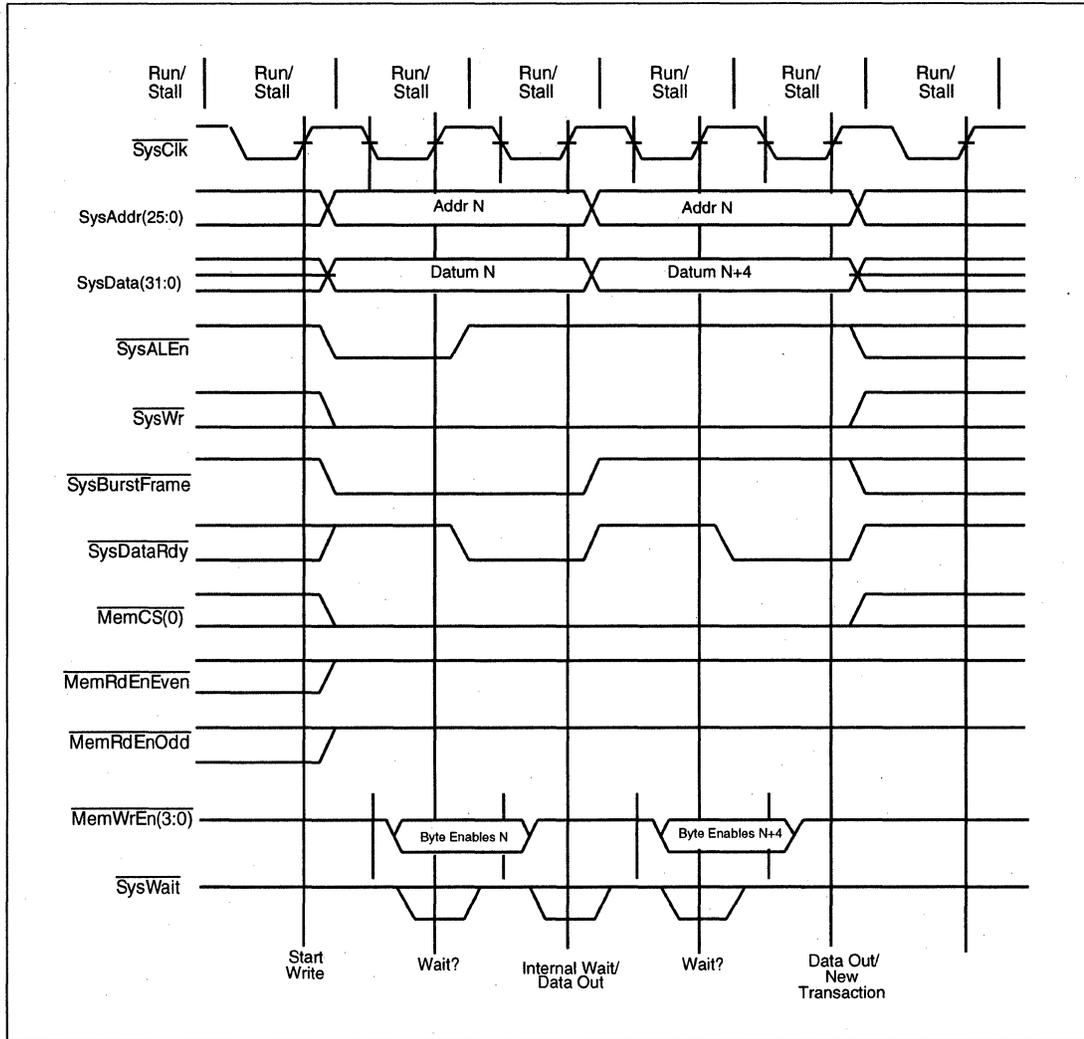


Figure 8.23 Multi-Datum Burst Write using Wait-State with WrDatum2Datum.

Basic PCMCIA-Type Memory Write with 0 Wait-States

Figure 8.24 illustrates a basic PCMCIA-Type Memory Controller write transaction. Each transaction begins with both SysALEn and SysBurstFrame asserting. At this time, SysWr asserts if hasn't done so already (from the previous transaction), and it is guaranteed that MemCS() will be in its de-asserted state. Assuming that there are no internally programmed StartRepeat wait-states, on the next clock cycle, SysBurstFrame de-asserts.

On the third cycle MemCS() asserts. On PCMCIA transactions the MemCS() pair is asserted according to which bytes are enabled and valid. If the even byte is valid, the even MemCS() will assert, while if the odd byte is valid, then the odd MemCS() will assert. If both bytes are valid, then both MemCS() signals in the pair will assert.

Assuming there are no internally programmed RdStart2Datum wait-states, SysDataRdy asserts to indicate that the data from the RISController is ready to be latched into the memory device. On the 4th clock cycle, MemCS() de-asserts, providing a means for the write data from the RISController to be latched into the memory device. On the next clock, the next transaction may begin.

During a PCMCIA-Type transaction, there is 1 clock of address setup time before MemCS() asserts. All signals are setup before MemCS(), which is being used as the write strobe, de-asserts.

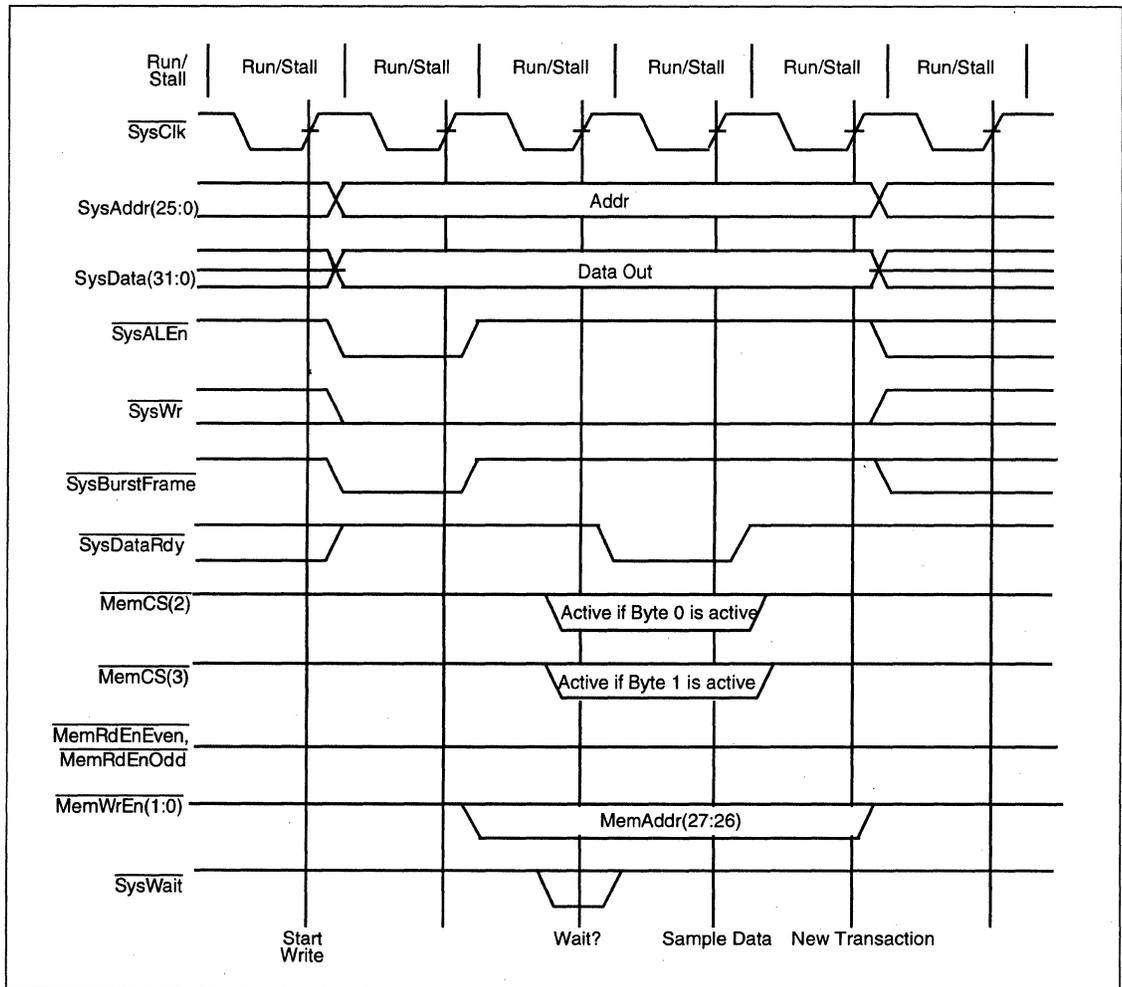


Figure 8.24 PCMCIA-Style Memory Write with 0 Wait-States

Interleaved-Type Transactions

The R36100 RISController's Memory Controller has the capability to interleave memory transactions such that 64-bits of memory are accessed at a time and then funneled into or out of the CPU in two 32-bit chunks. Thus on a burst read, the CPU begins a pair of 32-bit accesses at the same time. The first word is read in as normal, while the second word is externally latched. Then while the second word is read, the third and fourth words begin their accesses. When large memory arrays are used, interleaved systems speed up the overall transaction time of burst reads and do not add cost to the system, since transceivers are typically needed to isolate the multiple memory banks anyway. The R36100 supports a variety of different data transceiver options, as shown here.

Interleaved Read using FCT260-Type Field

Figure 8.25 illustrates an interleaved 4-word burst read using the FCT260-Type. The odd Read Enable, MemRdEnOdd is used as the Path Select and the Odd Latch Enable. MemRdEnEven changes its functionality in that it asserts for both the even and odd Data Sampling periods. As with the non-interleaved types, various throttled wait-state options are available via the internal Wait-State Generator including optimal Burst Ack placement. It is implied that the odd words are always returned 1 clock after the even words.

Figure 8.26 shows a single datum access to an "even" bank of an interleaved memory system using FCT260 transceivers. Note that the timing of this access is identical with the timing of the first word of a 4-word access.

Figure 8.27 shows the analogous access to the "odd" bank of an interleaved memory system, using FCT260 transceivers. In this figure, the timing is identical with the timing of the access of the *second* word of a 4-word access; however, the first word is not actually returned to the CPU.

Due to a limitation on the number of transceiver control signals, there is a performance difference between even and odd single word accesses. However, note that this should not affect system performance in the following areas:

- single word accesses will occur for uncached loads, and uncached instructions. These are typically not time critical programs or data.
- for cached data, these accesses will only occur if the data block refill option is selected to "one word" rather than "four words". However, this selection is unlikely for an interleaved memory, which dramatically mitigates the time required for 4-word accesses, and thus is expected to use 4-word accesses on data cache misses.
- cached instruction misses are always satisfied using 4-word read accesses.

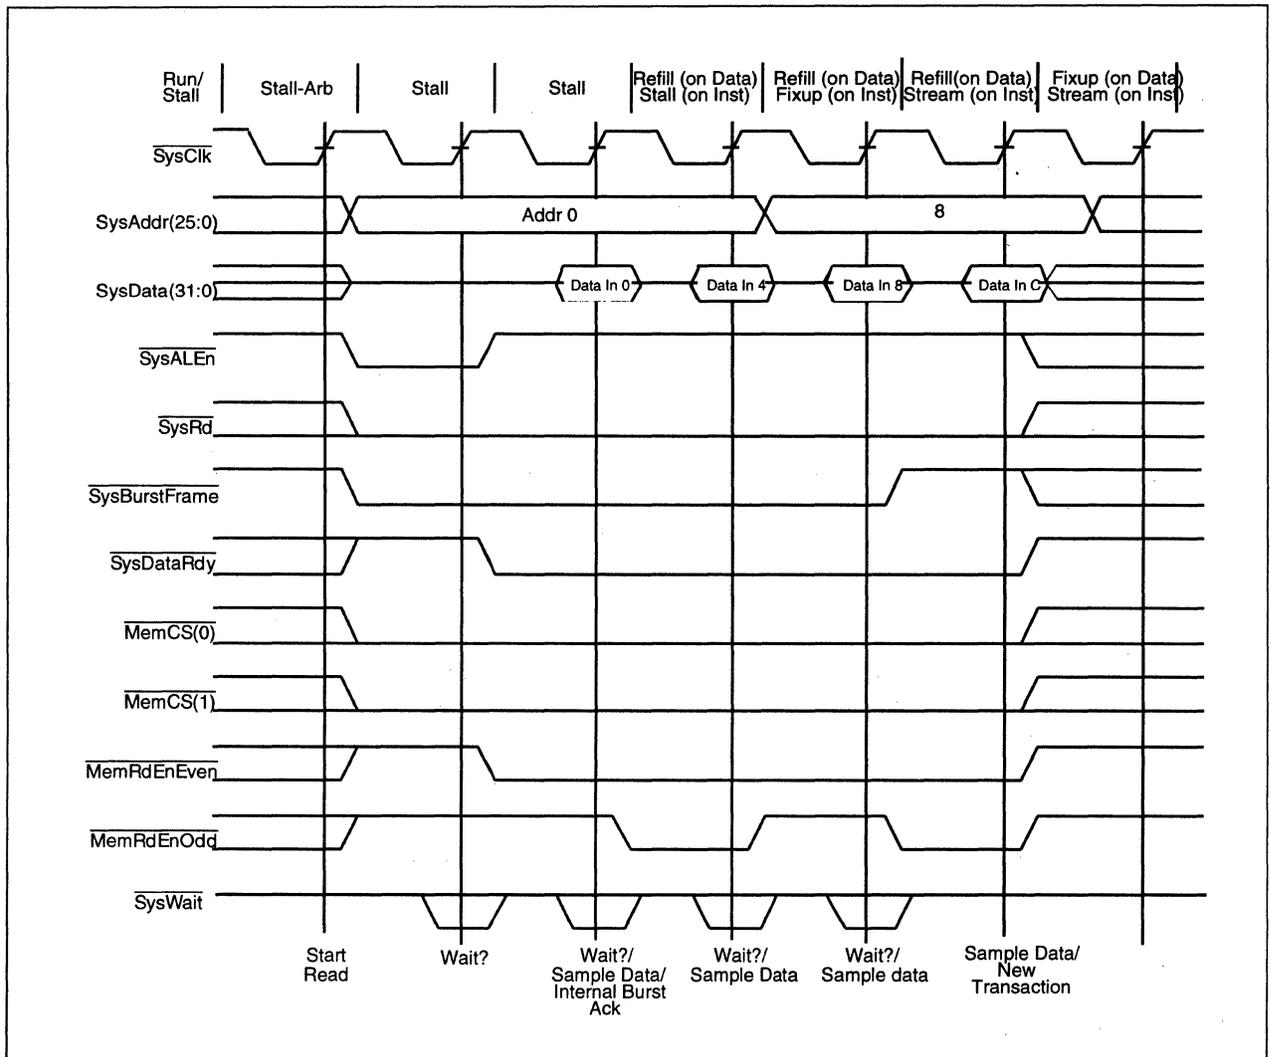


Figure 8.25 Interleaved Read using FCT260-Type Field.

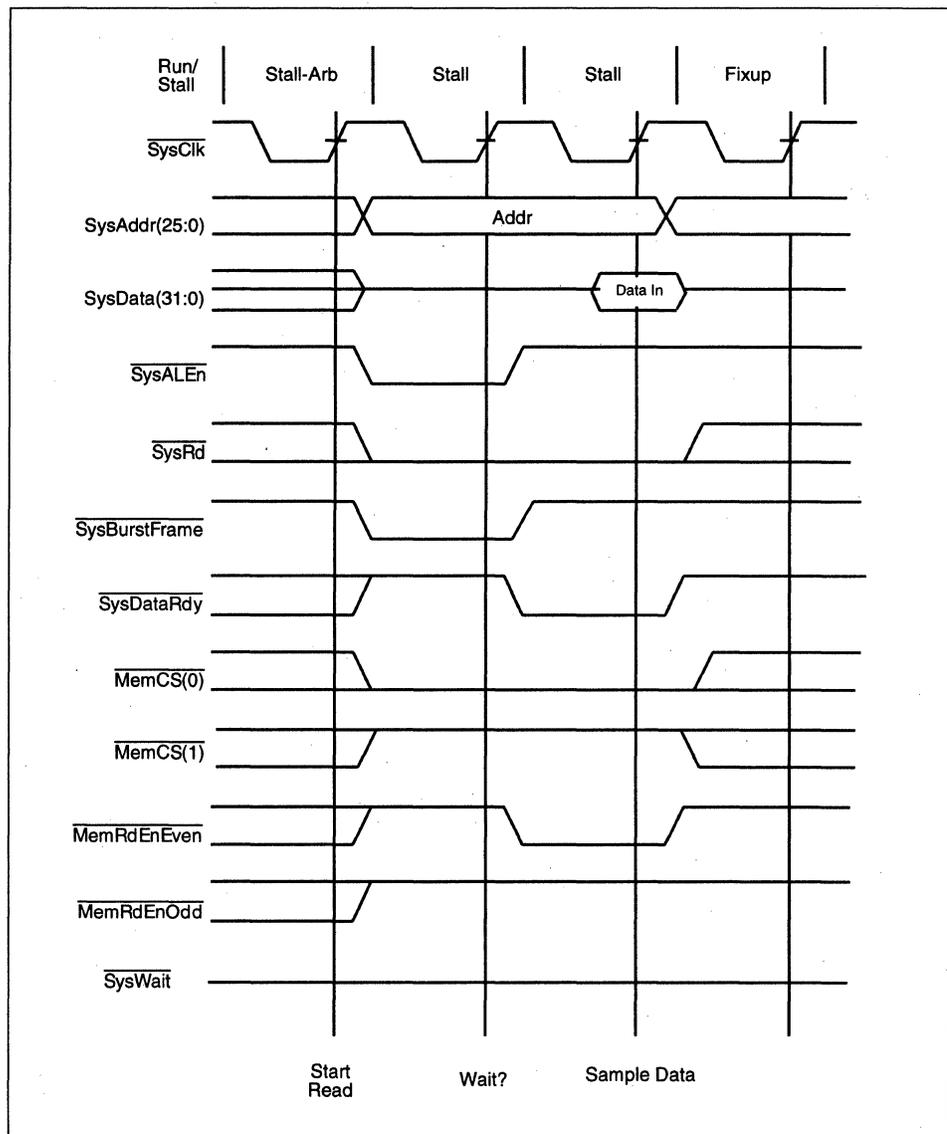


Figure 8.26 Interleaved "Even" Read of FCT260-Type Memory

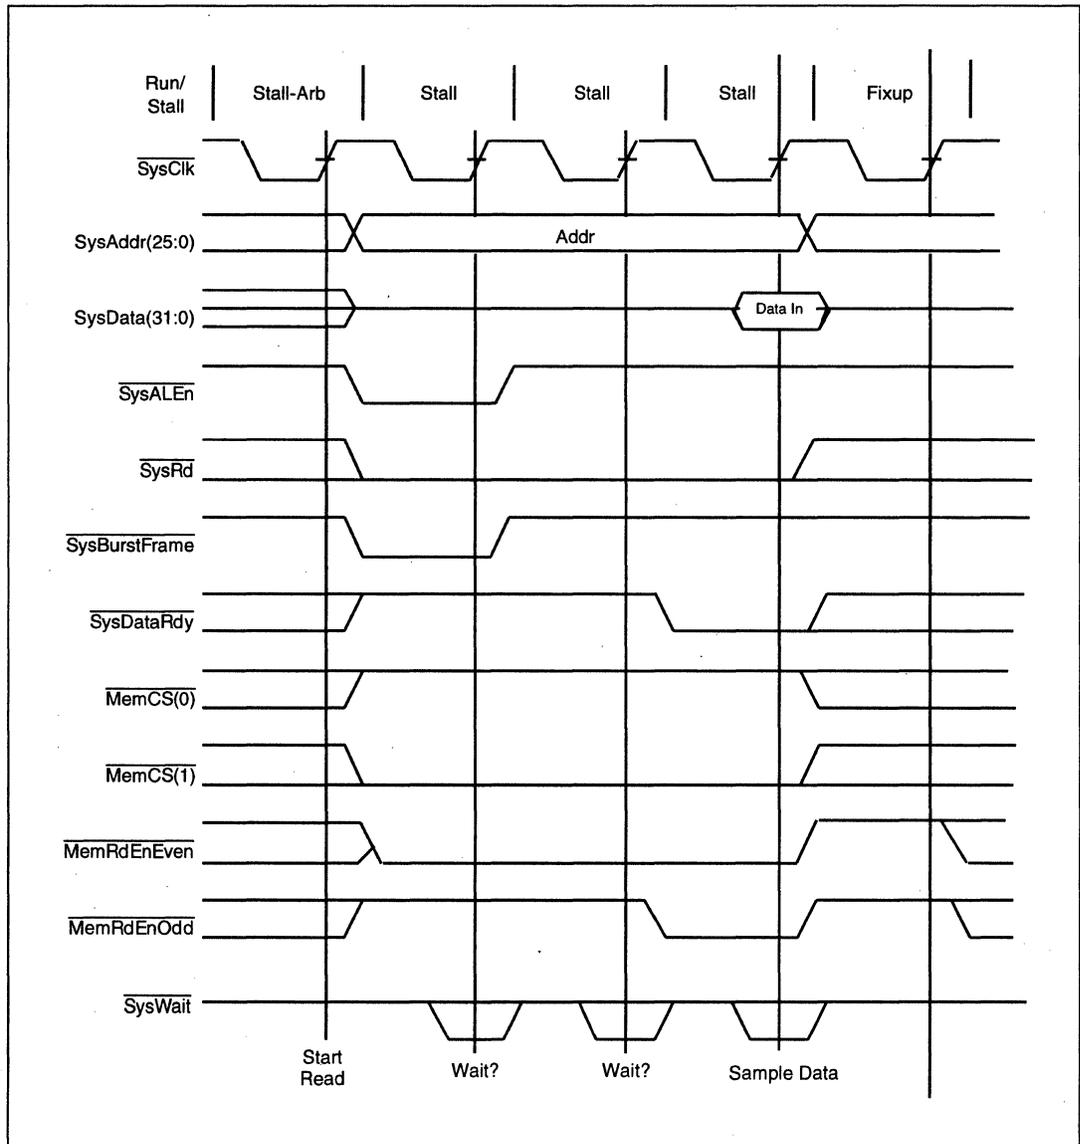


Figure 8.27 Interleaved "Odd" Read of FCT260-Type Memory

Interleaved Read using FCT245-Type Field

Figure 8.28 illustrates an interleaved 4-word burst read, using the FCT245-Type. The even and odd Read Enables, MemRdEnEven and MemRdEnOdd, change their functionality and become active for both reads and writes. This allows their use for the output enable pins of their respective FCT245 banks. Also, because FCT245s do not latch data, the doubleword increment from 0x0 to 0x8 is done after the second word is completely read and sampled by the R36100. The Datum2Datum delay is between pairs of datum.

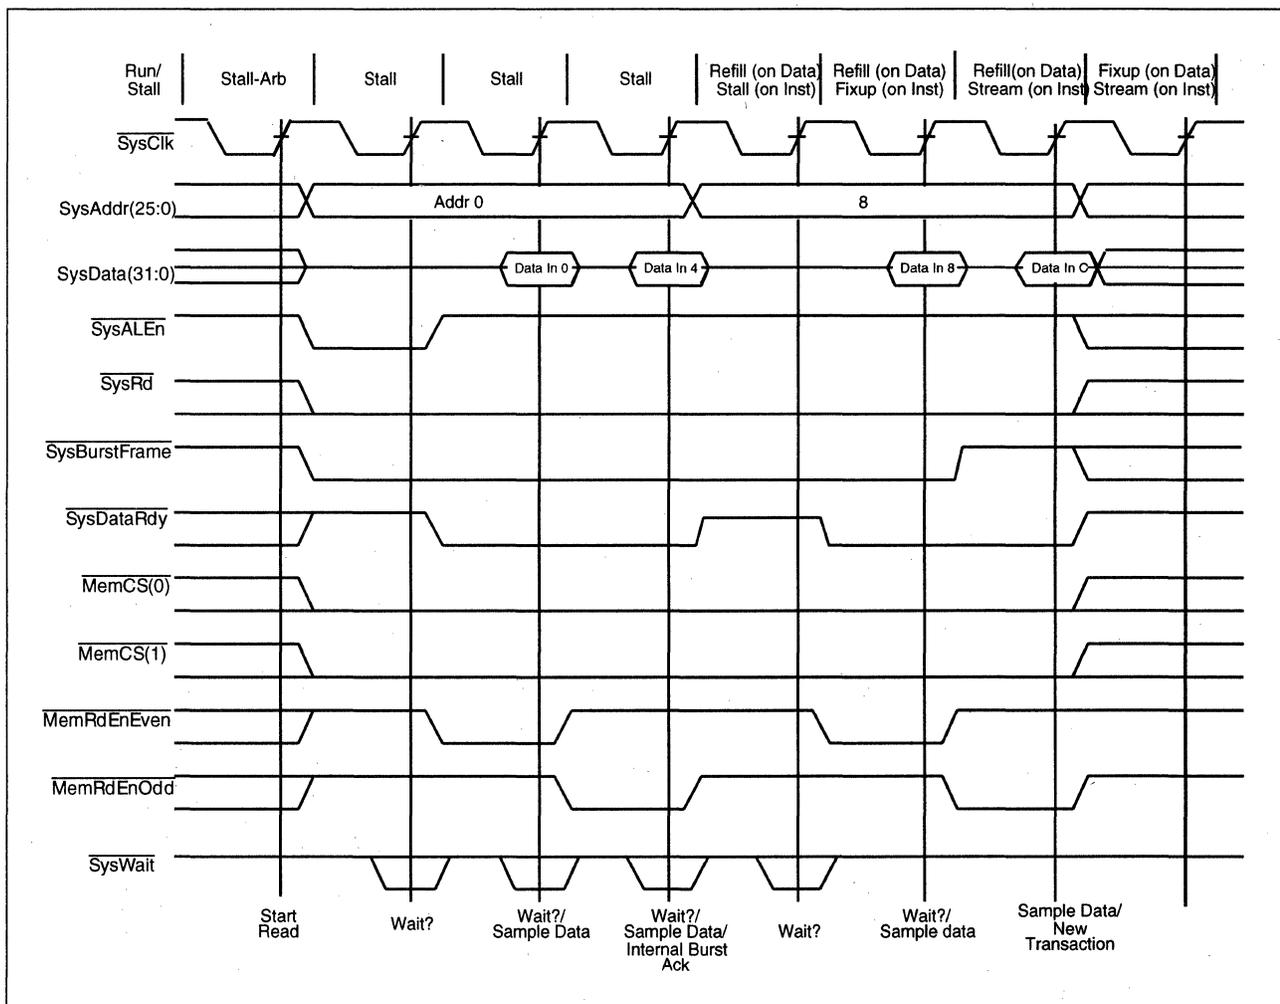


Figure 8.28 Interleaved Read using FCT245-Type Field.

Interleaved Read using FCT543-Type Field

Figure 8.29 illustrates an interleaved 4-word burst read using the FCT543-Type. The read and write output enables match the functionality of memory chip read and write output enables.

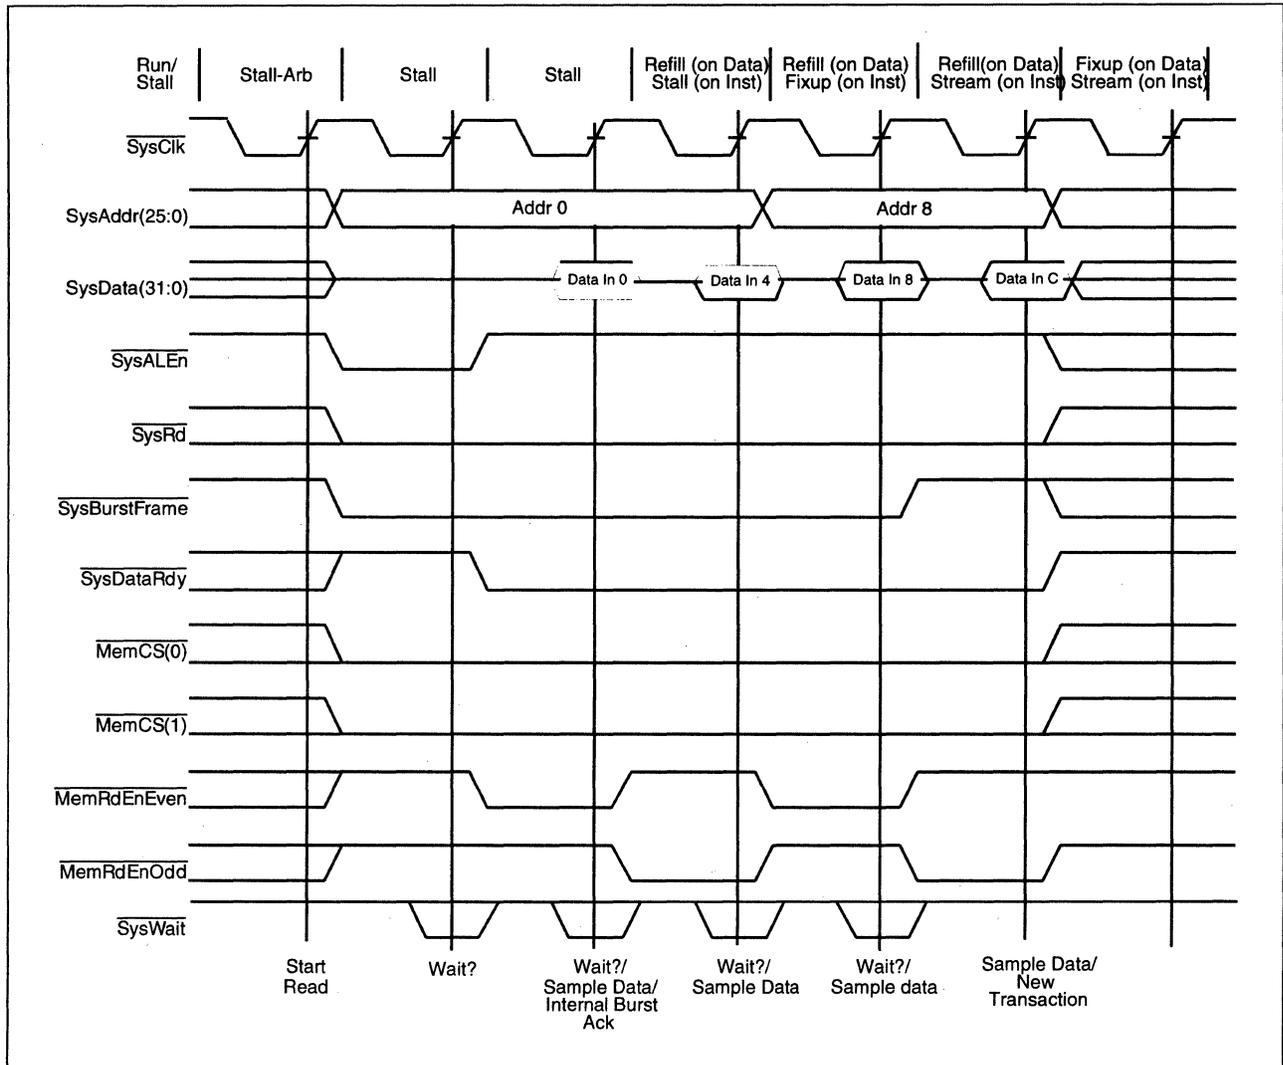


Figure 8.29 Interleaved Read using FCT543-Type Field.

The following Figure 8.30 shows a single datum access to an “even” bank of an interleaved memory system using 'FCT543 transceivers. Note that the timing of this access is identical with the timing of the first word of a 4-word access.

Figure 8.31 shows the analogous access to the “odd” bank of an interleaved memory system, using 'FCT543 transceivers. In this figure, the timing is identical with the timing of the access of the *second* word of a 4-word access; however, the first word is not actually returned to the CPU.

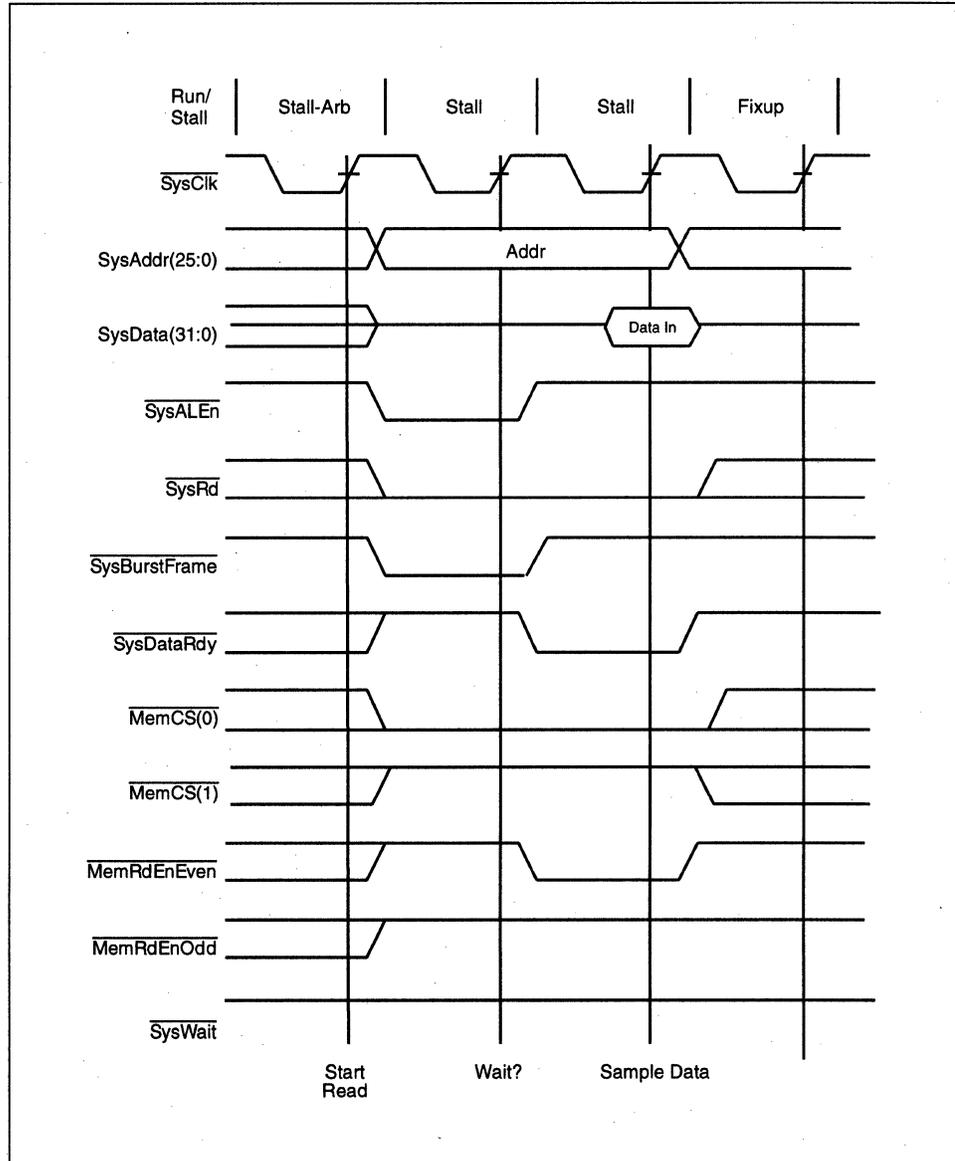


Figure 8.30 "Even" Read of FCT543-Type Memory

Due to a limitation on the number of transceiver control signals, there is a performance difference between even and odd single word accesses. Note, however, that this should not affect system performance of the following:

- single word accesses will occur for uncached loads, and uncached instructions. These are typically not time critical programs or data.
- for cached data, these accesses will only occur if the data block refill option is selected to "one word" rather than "four words". However, this selection is unlikely for an interleaved memory, which dramatically mitigates the time required for 4-word accesses, and thus is expected to use 4-word accesses on data cache misses.
- cached instruction misses are always satisfied using 4-word read accesses.

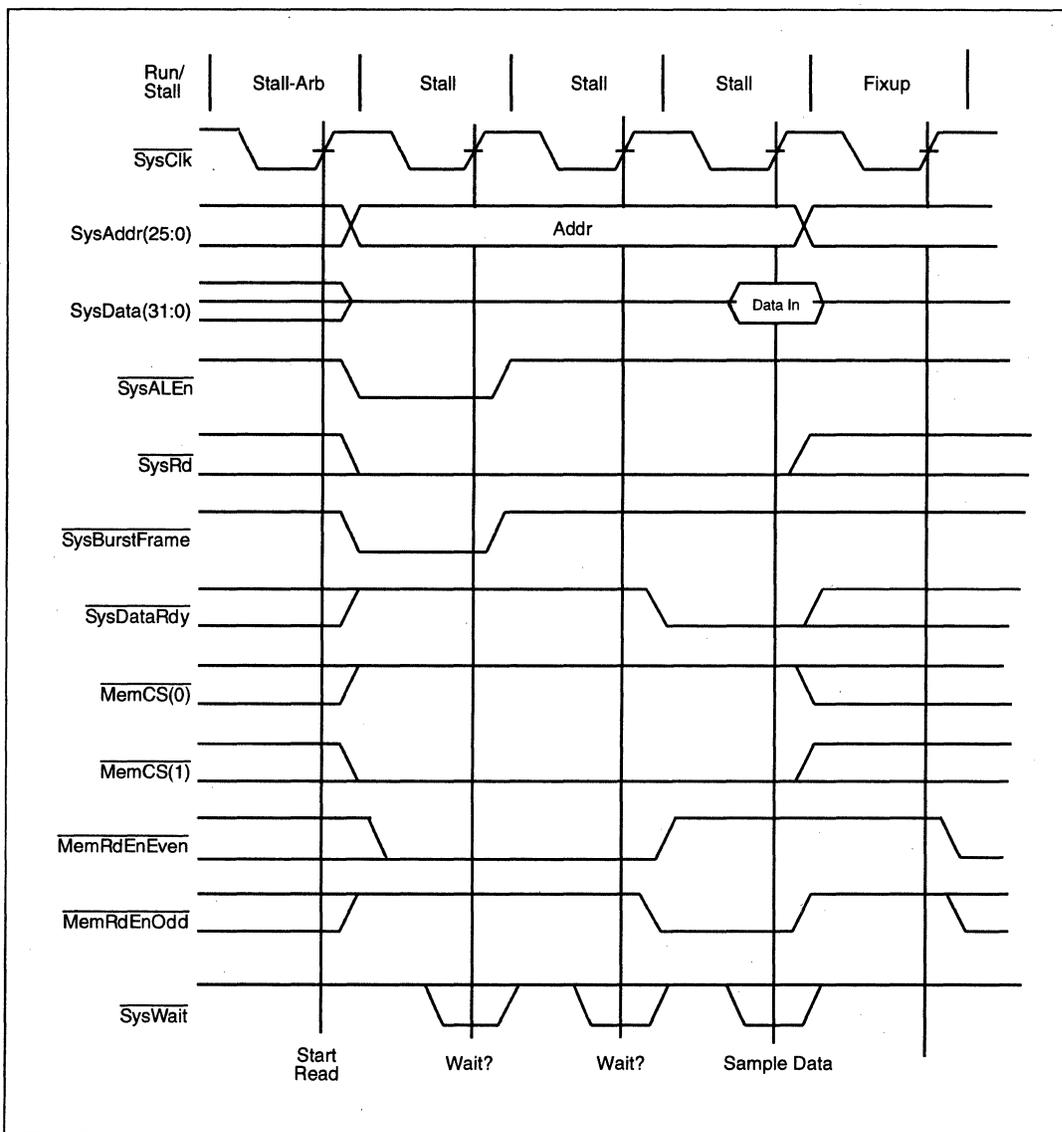


Figure 8.31 "Odd" Read of FCT543-Type Memory

Interleaved Writes

Figure 8.32 and Figure 8.33 illustrate interleaved writes for the 260-Type, the 543-Type, and the 245-Type, respectively. Because the byte enables, $\overline{\text{MemWrEn}}(3:0)$, are not duplicated for even and odd cases, burst writes do not occur any quicker than they do for non-interleaved cases. Any subsequent words are delayed by the $\overline{\text{MemWrEn}}(3:0)$ de-asserting for 1 clock (or 2 clocks if the Start Repeat Field is set). The 245-Type is different from the others in that $\overline{\text{MemRdEnEven}}$ or $\overline{\text{MemRdEnOdd}}$ is asserted for the 245's Output Enable pin.

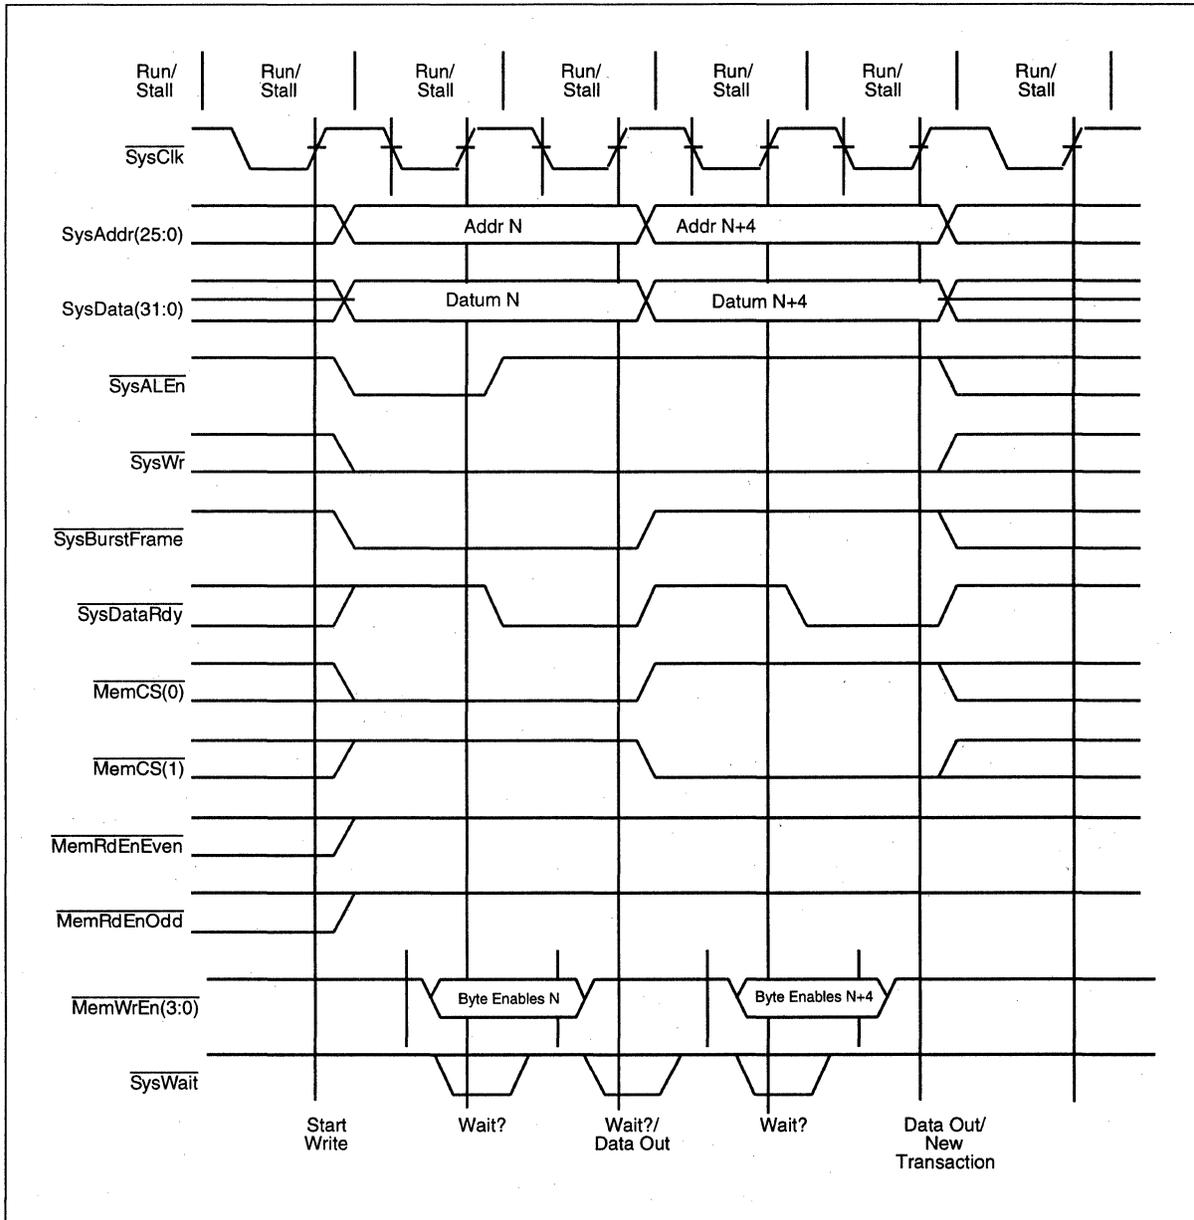


Figure 8.32 Interleaved Write using FCT260-Type and FCT543-Type Fields.

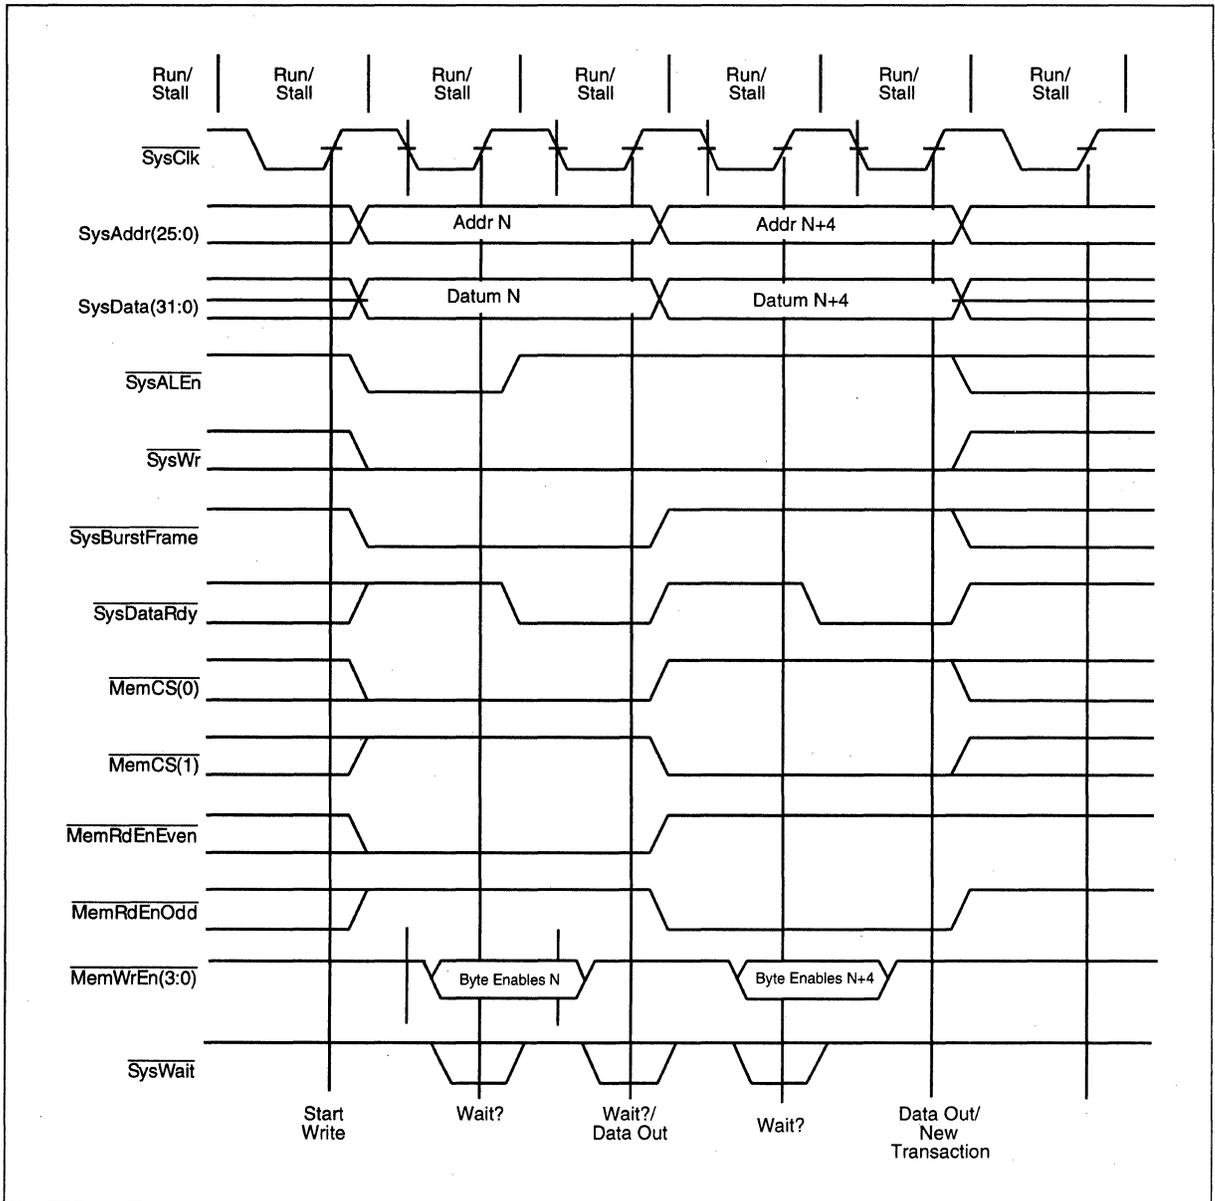


Figure 8.33 Interleaved Write using FCT245-Type Field.

System Examples

32-bit SRAM/ROM Directly Connected or using 543 Transceivers

Figure 8.34 shows a typical 32-bit SRAM memory system. The example is also applicable to Flash Memory systems and with the elimination of the write-related signals, to ROM systems. In small systems, the SRAM/ROM can be attached directly to the SysAddr and SysData buses. In larger systems, FCT543 transceivers can be added between the memory bank and the SysData bus. Also in large systems, the SysAddr bus can also be buffered using FCT244 buffers. Note that all even MemCS(0)s that use transceivers must be placed behind the same set of transceivers unless external decoding is done for each MemRdEnEven OR MemCS(0).

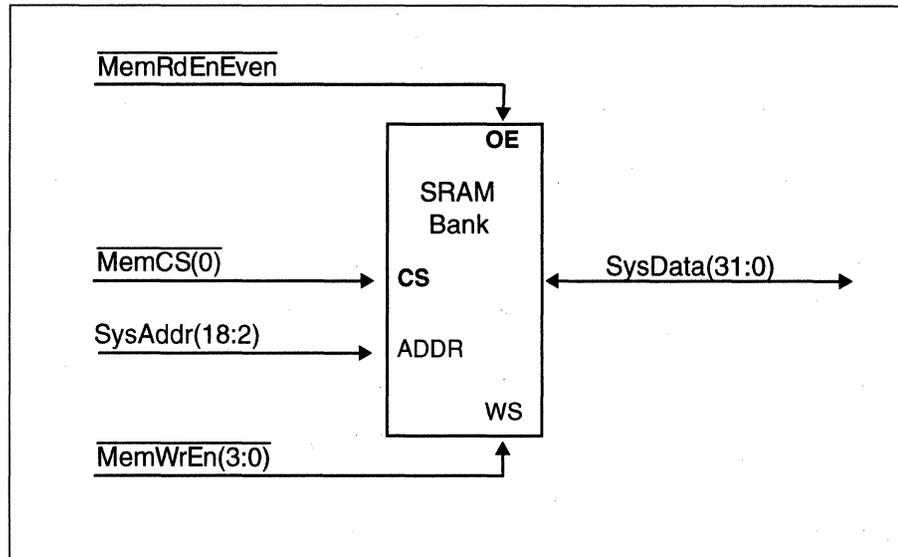


Figure 8.34 32-bit SRAM System.

Figure 8.35 shows a typical 32-bit SRAM memory system using an odd MemCS(1) line. Note that all odd MemCS(1)s that use transceivers must be behind the same set of transceivers unless external decoding is done for each MemRdEnOdd OR MemCS(1).

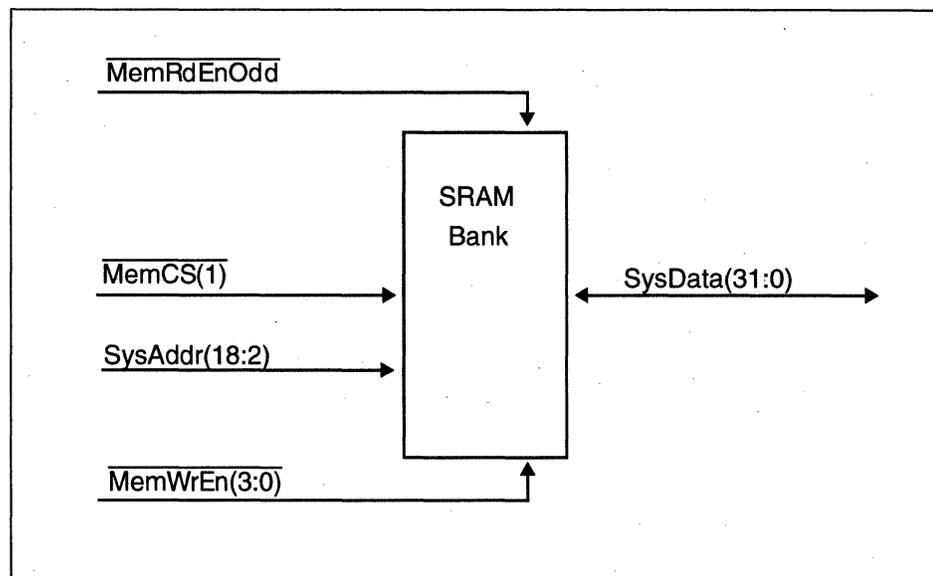


Figure 8.35 32-bit SRAM System using an Odd Chip Select.

32-bit SRAM using 245 Transceivers

Figure 8.36 shows a typical 32-bit SRAM memory system using FCT245 transceivers. This example is also applicable to Flash Memory systems and with the elimination of the write-related signals and the substitution of FCT244 buffers to ROM systems. The functionality of $\overline{\text{MemRdEnEven}}$ or $\overline{\text{MemRdEnOdd}}$ is modified so that the read enables will also assert on writes. The use of a 245-Type chip select requires that all the odd or all the even memory chip selects also be of the same 245-Type type.

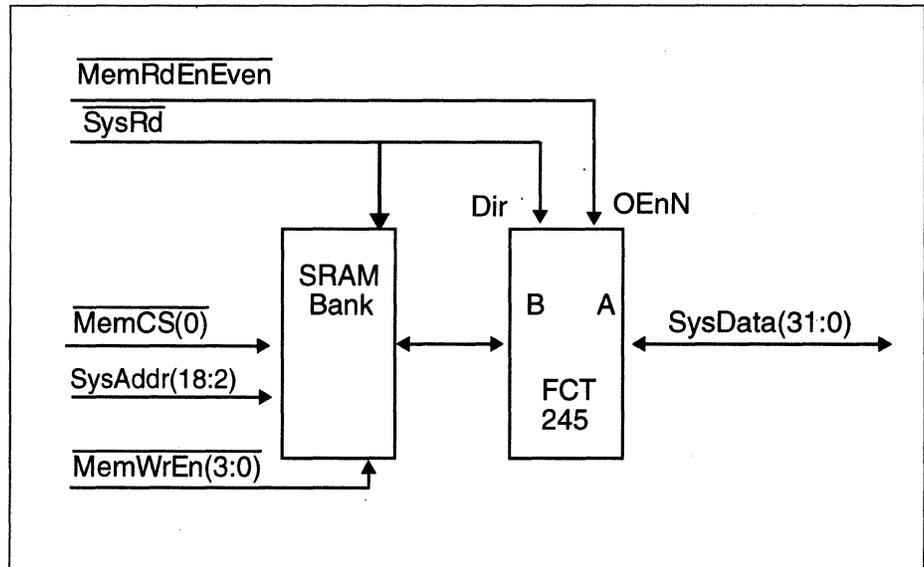


Figure 8.36 32-bit SRAM System using FCT245-Type.

Interleaved SRAM/ROM using 245 Transceivers

In Figure 8.37 an interleaved system using the Interleaved-FCT245-Type is shown. FCT245s are 16-bit (or 8-bit) bidirectional transceivers. Four (or eight) are required per bank. Because FCT245s do not latch the data of the odd words, the generation of the next double word address (0x8) for words three and four is delayed. MemRdEnEven and MemRdEnOdd change their functionality in that if their bank is selected, as they assert on *both reads and writes* such that they can be connected to their bank's 245 Output Enable pin.

Note: The R36100 cannot have a stand-alone 32-bit odd bank, because it must be interleaved with an even bank partner. For the FCT245, interleaved ROMs must be booted with 32-bit port size and load the initial first few instructions from every even word, up until the boot code initializes the memory controller registers

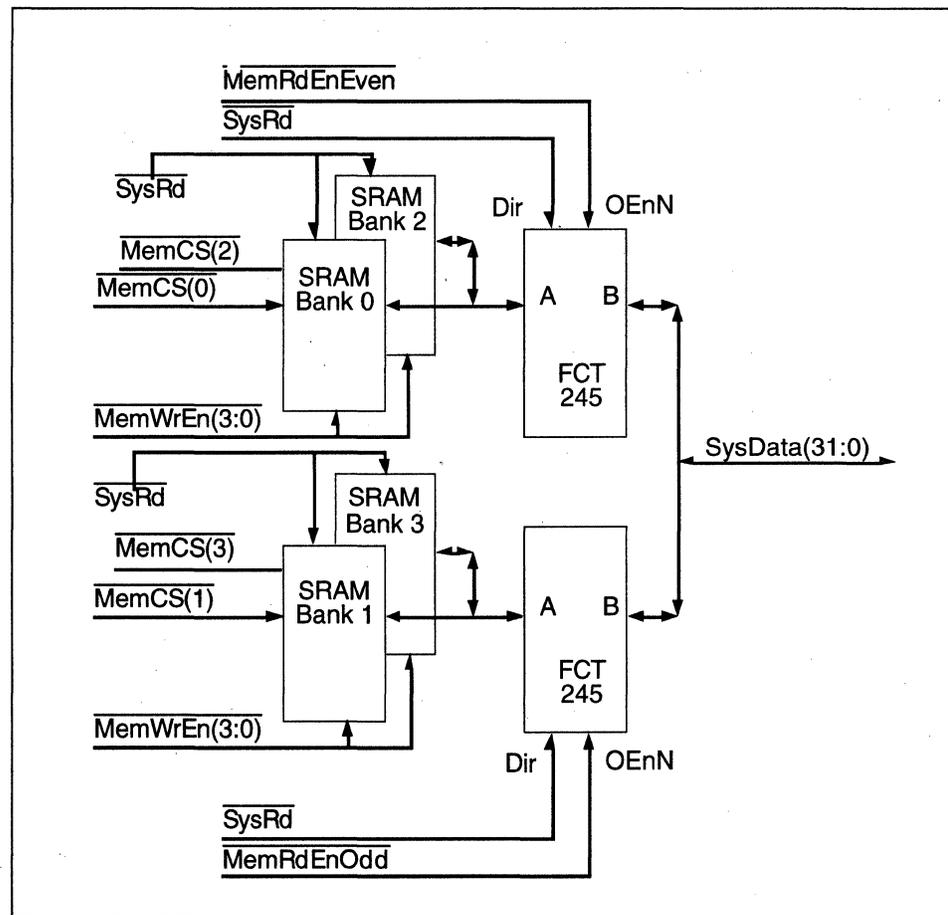


Figure 8.37 Interleaved FCT245-Type System.

Interleaved SRAM/ROM using 260 Multiplexors

In Figure 8.38 an interleaved system using the Interleaved-FCT260-Type is shown. FCT260s are 12-bit bidirectional multiplexors. Thus three FCT260s are required per bank pair. $\overline{\text{MemRdEnOdd}}$ is used for both the Path Selection and the Odd Bank Latch Enable. $\overline{\text{MemRdEnEven}}$ changes its functionality in that it asserts for both the even and odd Data Sampling periods. The LSB address lines change on a doubleword boundary from 0x0 to 0x8 at the same time word 2 is being read.

Note that during an even (odd) bank read access, the chip select (for example, CS0 or CS2) is used to keep the two banks from contending on the data bus. In this example, SysWr is used to control the data path outputs on writes.

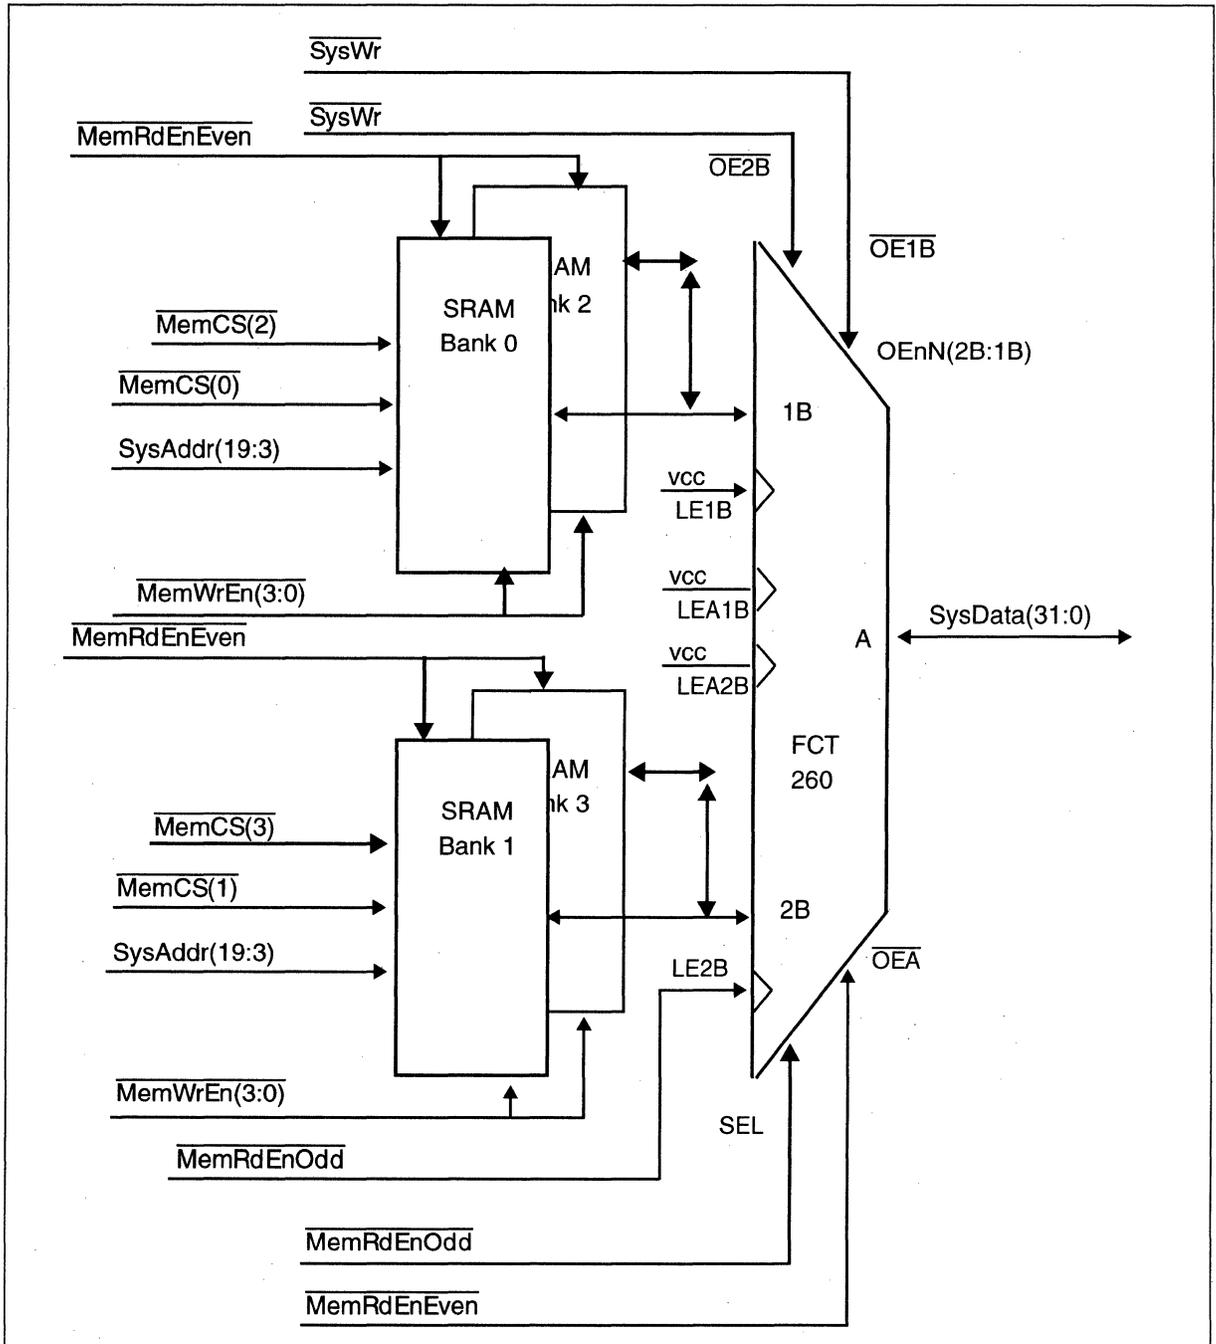


Figure 8.38 Interleaved FCT260-Type System.

Interleaved SRAM/ROM using FCT543 Transceivers

Figure 8.39 shows an interleaved system using the Interleaved-FCT543-Type. FCT543s are 16-bit bidirectional registers. Thus, four FCT543s are required per bank pair. MemRdEnOdd is used for both the Path Selection and the Odd Bank Latch Enable. MemRdEnEven changes its functionality in that it asserts for both the even and odd Data Sampling periods. The LSB address lines change on a doubleword boundary from 0x0 to 0x8 at the same time that word 2 is being read.

In systems that contain memories with fast output disable times, only two FCT543s may be required by letting MemRdEnEven and MemRdEnOdd control the outputs of the two memory banks behind a common transceiver. Such an approach usually implies that the two memory banks be of the same type and that the type specify output disable times less than output enable times.

Note: The R36100 cannot have a stand-alone 32-bit odd bank, because it must be interleaved with an even bank partner. For the FCT543, interleaved ROMs must be booted with 32-bit port size and load the initial first few instructions from every even word, up until the boot code initializes the memory controller registers

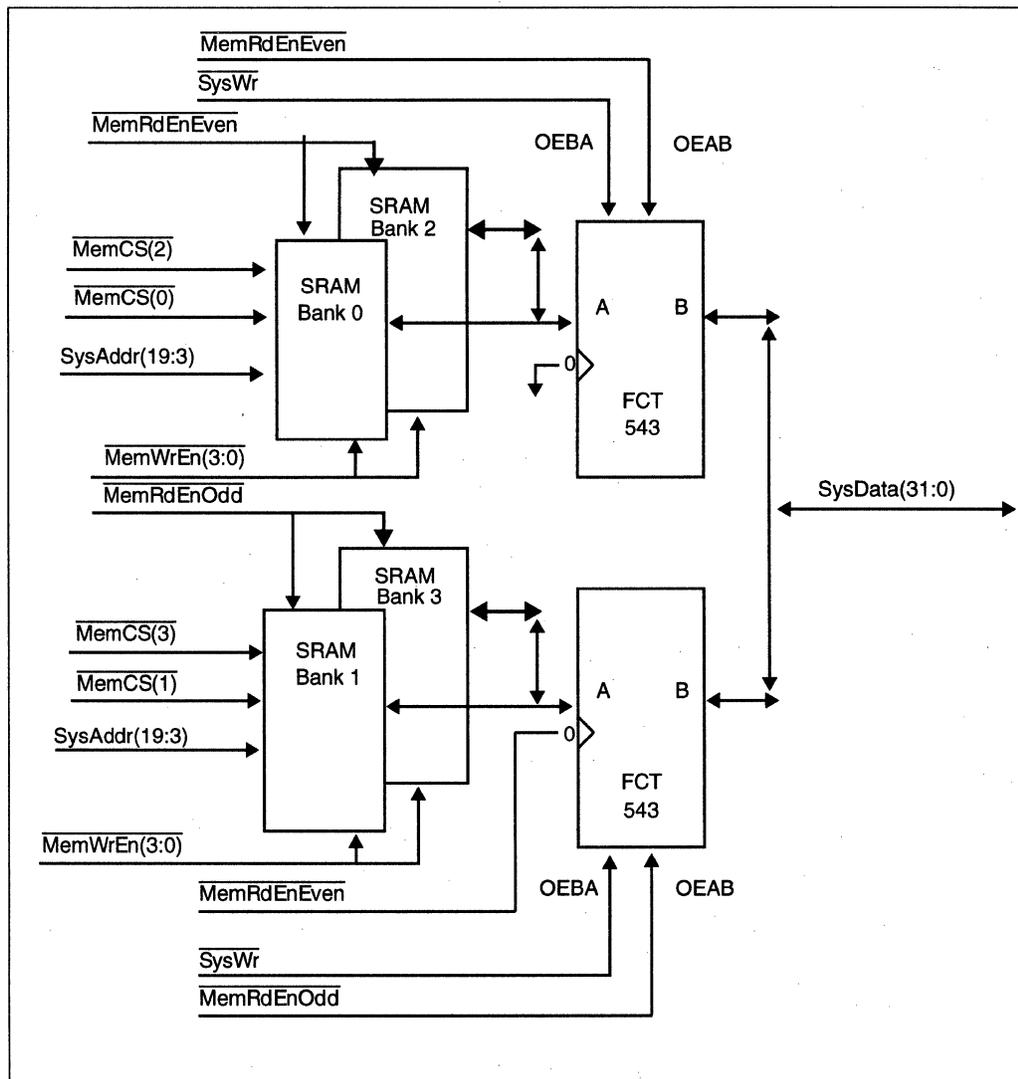


Figure 8.39 Interleaved FCT543-Type System

16-bit SRAM/ROM

Figure 8.40 and Figure 8.41 show a typical 16-bit SRAM memory system. These examples are also applicable to Flash Memory systems and, with the elimination of the write-related signals, to ROM systems. Since a 16-bit system is a smaller system, the SRAM/ROM can be attached directly to the SysAddr and SysData buses. In larger systems, FCT543 transceivers can be added between the memory bank and the SysData bus. Also in large systems, the SysAddr bus can be buffered using FCT244 buffers.

Note: With 16-bit systems it is imperative that the correct data line connections are made. Big Endian systems must attach SysData(31:16) and Little Endian systems must attach SysData(15:0).

Hooking up the correct data lines insures that byte gathering can occur on word accesses and is required for boot-proms to execute instructions. SRAM systems that require later expansion to 32-bits must externally multiplex a MSB address line with SysAddr(1).

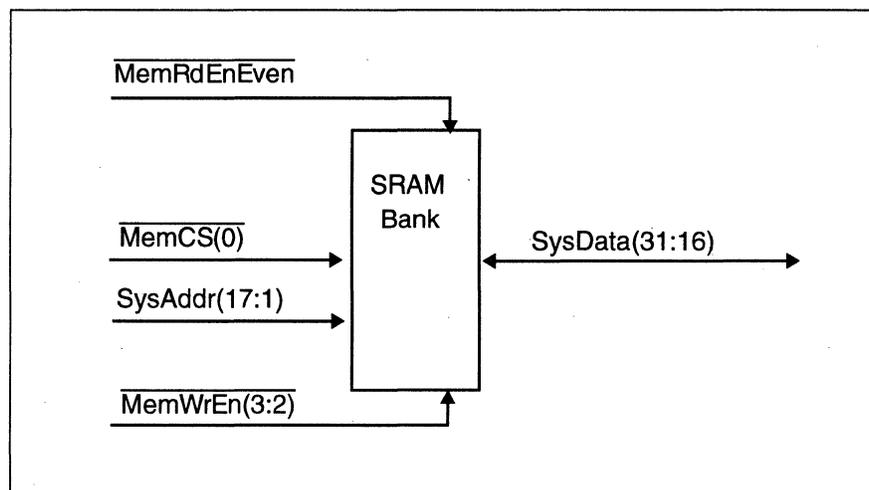


Figure 8.40 16-bit Big Endian SRAM System.

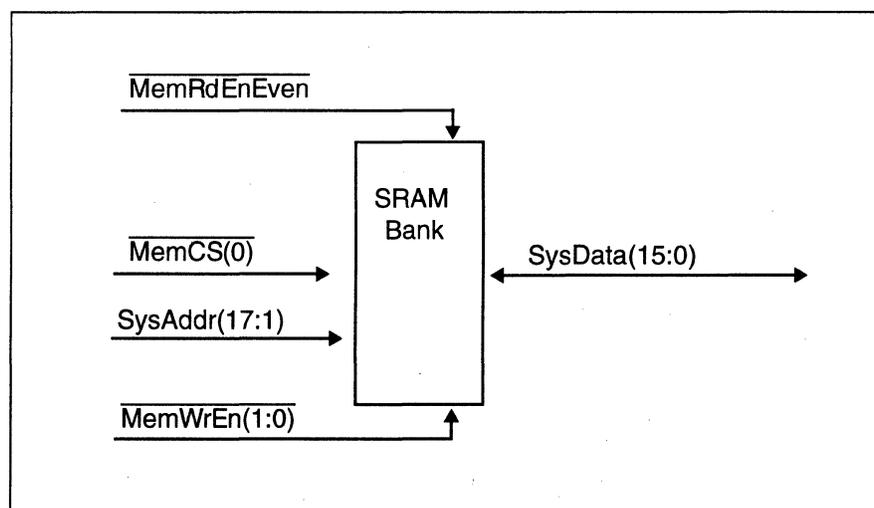


Figure 8.41 16-bit Little Endian SRAM System.

8-bit SRAM/ROM

Figure 8.42 and Figure 8.43 show two typical 8-bit SRAM memory systems. These examples are also applicable to Flash Memory systems and, with the elimination of the write-related signals, to ROM systems. Since an 8-bit system is a smaller system, the SRAM/ROM can be attached directly to the SysAddr and SysData buses. In larger systems, FCT543 transceivers can be added between the memory bank and the SysData bus. Also in large systems, the SysAddr bus can be buffered using FCT244 buffers.

Note: With 8-bit systems it is imperative that the correct data line connections are made. Big Endian systems must attach SysData(31:24) and Little Endian systems must attach SysData(7:0).

Hooking up the correct data lines insures that byte gathering can occur on word accesses, and it is required for boot-proms to execute instructions. SRAM systems that require later expansion to 32-bits must externally multiplex a MSB address line with SysAddr(0).

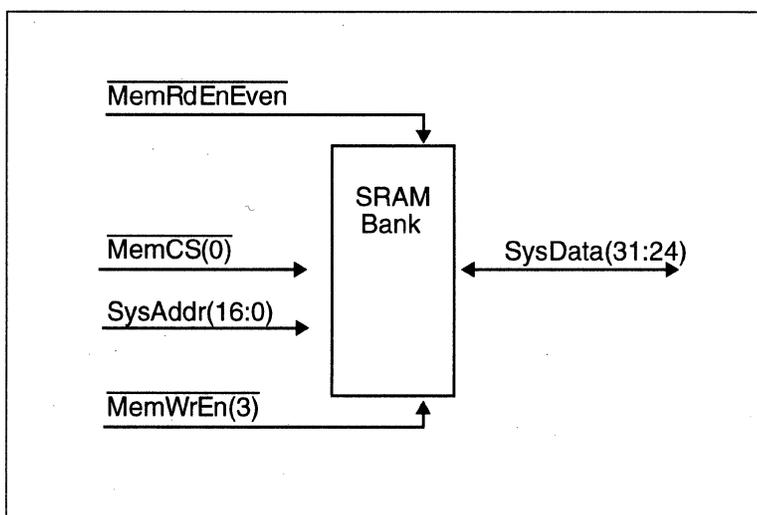


Figure 8.42 8-bit Big Endian SRAM System.

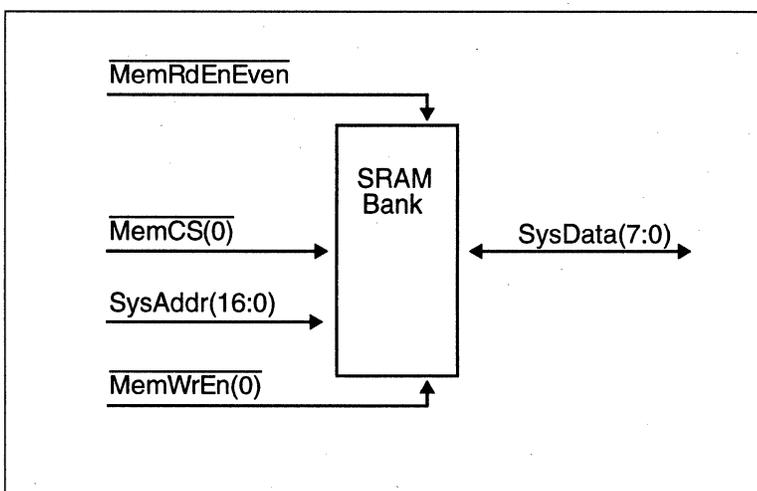


Figure 8.43 8-bit Little Endian SRAM System.

Dual-Port-Type

Dual-Port memory systems usually have a BusyN pin that indicates both ports of a memory location are being accessed simultaneously. The port that is accessed second receives the BusyN signal, indicating it must wait until the port that was accessed first is finished. Such Dual-Port systems can use the SysWait signal.

This allows a full Dual-Port memory access cycle when BusyN deasserts. Using this type, if the Dual-Port memory glitches BusyN, (for instance, if the addresses match before MemCS is asserted, but don't afterwards) then the Dual-Port access will not be optimal in the sense that additional wait-states will be injected; however, operation will be correct. If increased optimization is required, the system designer can add external circuitry to gate BusyN with the beginning of MemCS so that it is ignored until BusyN is valid.

PCMCIA-Style Application

Multiple cards can be externally decoded using SysWrEn(1:0) programmed as MemAddr(27:26). Table 8.16 lists the PCMCIA and R36100 functional equivalents.

PCMCIA	R36100 Function
A	SysAddr(25:0).
REG	Discrete PIO output pin.
~CE(2:1)	MemIoCS() pair program defaulted to PCMCIA-Memory Style and then to PCMCIA-IO-Style when necessary.
~IORD	IoRd.
~IOWR	IoWr.
~OE	~(~MemIoCS) & ~ SysRd & ~ SysDataRdy)
~WE	MemWrEnEven.
D	SysData(15:0).
~WAIT	SysWait
~INPACK	Ignored.
~IOIS16	Ignored.
RESET	System dependent or PIO output pin.
~IREQ	ExcInt() pin.
CD(2:1)	Two PIO input pins.
BVD(2:1)	Two PIO input pins or ignore.
WP	System dependent.

Table 8.16 PCMCIA and R36100 Functional Equivalents.

PCI-Style

The R36100 RISController does not directly support the PCI Bus specification. The intention is to allow PCI-style peripheral chips to be used, most of which are tolerant of slight variations on the PCI specification.

Thus certain liberties, such as lower drive and slightly looser AC setup and hold specifications, may have to be accommodated; certain features—such as device identification, parity, and various system status protocols—are ignored. It is assumed that the PCI-style peripheral control registers are loaded and read with slave accesses generated from the RISController and that the peripheral uses an External DMA Channel (see Chapter 11) to read and write data via a RISController Bus Controller.

PCI-Style changes the $\overline{\text{MemWrEn}}(3:0)$ pins to PCI-Style pins using the MSB Wait State Control Register BEn Field. $\overline{\text{MemWrEn}}(3:0)$ becomes active for both reads and writes. The $\overline{\text{MemWrEn}}$ pins always return high at the end of the bus transaction. In addition, the sense of the $\overline{\text{SysWait}}$ is inverted using the MSB Wait State Control Register InvWait Field. Note that PCI has to have address and data multiplexing done externally, typically with three FCT260 bus exchangers. Table 8.17 lists the PCI and R36100 Functional Equivalents.

PCI	R36100 Function
FRAME*	Use $\overline{\text{SysBurstFrame}}$. Bursts must be aligned to the block size.
AD(31:0)	Use 3x FCT260/272 to multiplex $\overline{\text{SysAddr}}(25:0)$ and $\overline{\text{SysData}}(31:0)$ as selected by $\overline{\text{SysALEn}}$. " $\overline{\text{Sys}}\overline{\text{Addr}}(31:26)$ use pull-downs. PAL needed to generate the three OEN's.
C/BE*(3:0)	Use PAL to multiplex pull-ups, pull-downs, and Rd^* with $\overline{\text{MemWrEn}}(3:0)$ in $\overline{\text{SysBEn}}$ mode.
PAR	Ignore or use parity generating/checking transceiver instead of FCT260/272.
IRDY*	$\overline{\text{SysDataRdy}}$
TRDY*	$\overline{\text{SysWait}}$ in InvWait mode on Slave Reads. Data must be latched with FCT260/272 for 1 clock.
DEVSEL*	Pull-Down; or FCT244 gated with $\overline{\text{DmaGnt}}()$.
STOP*	Ignore or use PAL (needed if lengthy linear bursts must be split into smaller/single transactions).
REQ*	$\overline{\text{DmaReqN}}()$
GNT*	$\overline{\text{DmaGntN}}()$

Table 8.17 PCI and R36100 Functional Equivalents.



Integrated Device Technology, Inc.

Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. One of the on-chip bus controllers is the “I/O Controller” as described in this chapter.

This chapter will provide an overview of the I/O Controller interface, a complete description of the signal pins and their timing, and how the interface relates to typical external hardware I/O devices and peripherals.

Because the I/O Bus Controller shares the Memory Controller (see Chapter 8) Registers and Chip Selects, the Register Description is not repeated here, except for the I/O specific “types” as described in the Type Field of each chip select's Control Register.

Features

- Controls support industry standard peripherals:
 - I-Type I/O support
 - M-Type I/O support
 - PCMCIA-Style I/O support
- Controls up to 8 banks of I/O

Note: Chip selects are shared with the Memory Controller.

- Each $\overline{\text{IOCS}}$ can be programmed to:
 - Individual chip selects
 - Combined PCMCIA-I/O-style pair-wise chip selects
- Each Bank has Programmable Base Address
- Each Bank Size programmable from 8KB - 64MB
- 8, 16, and 32-bit support
- Wait State Generator features:
 - Programmable time from start to end of each data access for each area
 - Programmable time options for reads and writes
 - Programmable time options for single word accesses
 - Internally generates the RdCEnN and AckN timing for all CPU accesses
 - A programmed value may be overridden by the $\overline{\text{SysWait}}$ input signal
- Direct control of data path transceivers supports various options:
 - Direct Bus Connection
 - FCT245 Bidirectional Transceiver
 - FCT543 Bidirectional Registered Transceiver

Block Diagram

Figure 9.1 is a functional block diagram of the Memory and I/O Controller. The main Memory and I/O Controller Control Signal State Machine is responsible for generating the basic I/O Control signals used to connect to external peripherals. These signals include chip selects, read enables/strobes, and write enables/strobes.

The I/O Controller works in cooperation with the Bus Interface Unit as described in Chapter 7. Thus the Control Signal State Machine sends and receives information from the BIU Controller for assistance with controlling the port width and controlling partial word reads and writes. The Control Signal State Machine also uses information stored in the software programmable I/O Controller Register Bank for example, to control I-Type versus M-Type accesses.

The Wait-State Generator takes care of sending and receiving information from the BIU Controller in order to control the sequencing and timing of reading and writing each individual datum. The number of wait-states is derived from the settings programmed into the Register Bank. Once the correct number of wait-states has been counted out, then the Wait-State Generator sets the appropriate internal BIU 'Acknowledge' signals. With the programmable Wait-State Generator it is possible to eliminate the external state machines that are traditionally used for this function.

The I/O Controller Decoder constantly monitors the Bus Interface Unit's address and data bus to see if either:

- The access is to the I/O Controller's Register Bank.
- The access is in one of the I/O Controller's Chip Select Areas that is responsible for controlling the bus transaction.

The Register Bank allows the software programmer access to the many different options of the I/O Controller. The chip select address ranges, the number of wait-states, the port-width of the chip select, and other similar options are programmed into the Register Bank as part of the software initialization sequence of the boot operating system.

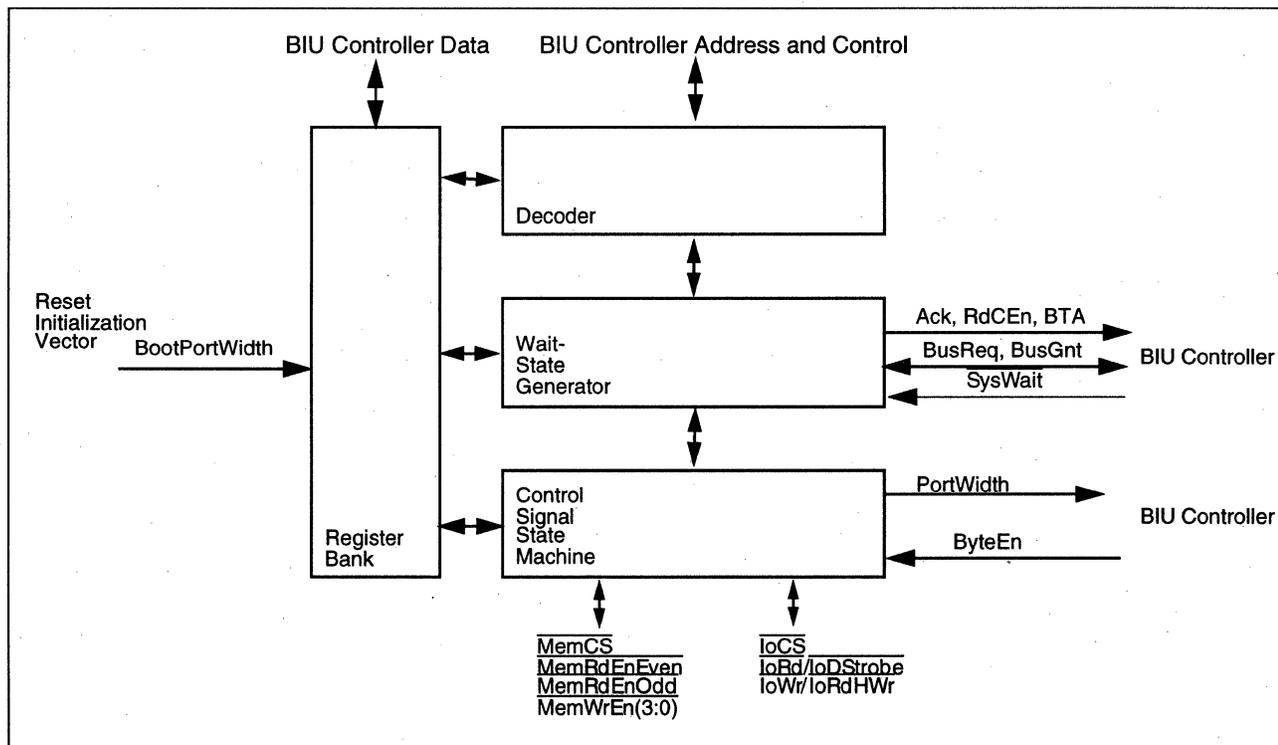


Figure 9.1 R36100 I/O Bus Controller Block Diagram

I/O Bus Controller Interface Signals

These external interface pins are typically attached directly from the R36100 RISController to external peripheral chips and their transceivers. Their descriptions are as follows:

$\overline{\text{IoCS}}(7:0)$ / Output

MemCS(7:0)

I/O Chip Select: The $\overline{\text{IoCS}}$ signals are active low outputs used to select one of the programmable I/O controller areas. Typically each external peripheral is attached to a $\overline{\text{MemCS}}$ signal such that the peripheral can be selected and turned on during an I/O transaction. When the address from the CPU or DMA Controller matches the I/O memory block corresponding to a particular $\overline{\text{IoCS}}$ signal, that $\overline{\text{IoCS}}$ asserts at the beginning of the next bus transaction and de-asserts at the end of that transaction.

$\overline{\text{IoCS}}$ signals are used individually for non-PCMCIA systems or in pairs (for example $\overline{\text{IoCS}}$ 0 & 1, 2 & 3, 4 & 5, or 6 & 7) for PCMCIA systems.

The boot PROM memory is assigned to $\overline{\text{Mem/IoCS}}(0)$ and if interleaved, $\overline{\text{Mem/IoCS}}(1)$.

The $\overline{\text{IoCS}}$ chip selects are selectable and shared between Memory and Mem-types (see Chapter 8).

$\overline{\text{IoRd}}$ / Output

IoDStrobe

This output is multiplexed depending on the I/O-Type selected in the Control Register of each I/O Space. This output only asserts during an I/O Controller bus transaction.

Input/Output Read: This active low output signal is used as a read enable strobe in conjunction with the write enable strobe, $\overline{\text{IoWr}}$. $\overline{\text{IoRd}}$ controls when the peripheral chip can drive the data signals back on to the main system data bus, SysData(). The timing of $\overline{\text{IoRd}}$ is such that SysAddr is stable before and after $\overline{\text{IoRd}}$ asserts.

Input/Output Data Strobe: This active low output signal is used as a data strobe for both reads and writes. $\overline{\text{IoDStrobe}}$ works in conjunction with the $\overline{\text{IoRdWr}}$ write status line. $\overline{\text{IoDStrobe}}$ controls when the data bus is valid. The de-asserting edge of $\overline{\text{IoDStrobe}}$ can be used to strobe the data into the peripheral on writes and indicates that the data has just been clocked into the CPU on reads. The timing of $\overline{\text{IoDStrobe}}$ is such that SysAddr is stable before and after $\overline{\text{IoDStrobe}}$ asserts. In some cases, $\overline{\text{IoCS}}(0)$ can be used instead of $\overline{\text{IoDStrobe}}$.

$\overline{\text{IoWr}}$ / Output

IoRdHWr

This output is multiplexed depending on the I/O-Type selected in the Control Register of each I/O Space. This output only asserts during an I/O Controller bus transaction.

Input/Output Write: This active low output signal is used as an I/O write strobe in conjunction with the I/O read enable strobe, $\overline{\text{IoRd}}$. $\overline{\text{IoWr}}$ controls when data is valid on the system data bus, SysData. When $\overline{\text{IoWr}}$ de-asserts, the peripheral can strobe the data into the chip. The timing of $\overline{\text{IoWr}}$ is such that SysAddr is stable before and after $\overline{\text{IoWr}}$ asserts.

Input/Output Read High and Write: This output is active high during I/O reads and active low during I/O writes. $\overline{\text{IoRdHWr}}$ is used as a read versus write status line. $\overline{\text{IoRdHWr}}$ works in conjunction with the $\overline{\text{IoDStrobe}}$. The timing of $\overline{\text{IoDStrobe}}$ is such that both $\overline{\text{IoRdHWr}}$ and SysAddr is stable before and after $\overline{\text{IoDStrobe}}$ asserts. Note that since this signal asserts on the same clock edge as $\overline{\text{IoCS}}()$, if setup is required to $\overline{\text{IoCS}}()$ as well, then $\overline{\text{SysWr}}$ (which asserts a clock earlier) can be substituted for $\overline{\text{IoRdHWr}}$.

**MemAddr(29:26) /
MemWrEn(3:0) /
MemByteEn(3:0)**

Output/(Input during DMA)

Memory Address Bus: During a PCMCIA Memory or I/O type access, the $\text{MemWrEn}(3:0)$ bus is instead driven with Physical Address bits. On the R36100 and R3051-base family memory map, virtual and physical addresses (29:0) are the same. An application using PCMCIA can, for example, use $\text{MemAddr}(27:26)$ to externally decode PCMCIA style chip selects into as many as four (256M/64M = 4) slots. In this mode (as in the other modes), the signals all return inactive high at the end of the bus transaction.

BIU Controller Signals

Many of the BIU Controller Signals are necessary to complete the I/O interface. These signals are listed here as a reminder. Information specific to the I/O Controller is given here and general information about the signal is given in Chapter 7, "System BIU Controller."

SysAddr

Output/Input

System Address Bus: SysAddr is an output bus when used with the I/O Controller. The MIPS architecture does not provide distinct memory and I/O spaces; thus MIPS I/O is considered to be "memory mapped I/O." A 32-bit peripheral connects to the word offset of the Least Significant Bits (LSB) of SysAddr . Thus such a 32-bit peripheral skips $\text{SysAddr}(1:0)$ and connects starting with $\text{SysAddr}(2)$ on up. A 16-bit peripheral may connect to the halfword offset of the LSBs of SysAddr starting with $\text{SysAddr}(1)$ on up. An 8-bit peripheral may connect the LSBs of SysAddr starting with $\text{SysAddr}(0)$ on up. 16-bit and 8-bit peripherals traditionally use the word offset on MIPS systems so that they can be addressed from either Big or Little endian data paths.

SysData

Output/Input

System Data Bus: A 32-bit peripheral connects the entire 32-bit SysData bus to its data pins or to its transceivers. A system implementation choice must be made on 16-bit and 8-bit peripherals.

In case 1, traditionally, MIPS systems have connected 16-bit and 8-bit peripherals with a word offset address. In such a case, a 16-bit device can be attached to $\text{SysData}(31:16)$ or $\text{SysData}(15:0)$ and accessed with either endianness by using a half-word offset of 0x2 for the opposite endianness. In such a case, an 8-bit device can be attached to $\text{SysData}(31:24)$ or $\text{SysData}(7:0)$ and accessed with either endianness by using a byte offset of 0x3 for the opposite endianness.

In case 2, 16-bit devices can use the halfword offset address, SysAddr(x:1) and 8-bit devices can use the byte address, SysAddr(x:0).

Note: In this case, 16-bit and 8-bit peripherals connect to particular data pins depending on whether the Endianness of the system is Big Endian or Little Endian.

Thus 16-bit peripherals use SysData(31:16) if they are Big Endian and SysData(15:0) if they are Little Endian. 8-bit peripherals use SysData(31:24) if they are Big Endian and SysData(7:0) if they are Little Endian. The User Mode Reverse Endianness Bit in the CPO Status Register has no effect on the connections to SysData, however, it strongly recommended that the Reverse Endianness Bit not be used to "correct" an endianness connection as it does not function in the kernel mode.

SysWait

Input

System Wait Negated: SysWait can be used by an external source to add wait-states to the I/O Controller. Since the I/O Controller itself has a Wait-State Generator, SysWait typically is not needed and can be pulled-up with a resistor. The most likely application of SysWait is for an asynchronous memory event such as a Dual-Port Memory Busy signal which can be used to attach to SysWait to delay the beginning of a I/O Transaction in the Wait Mode option of the Memory I/O Control Register 2. Please see Chapter 7 for a general description of SysWait.

Overview of the I/O Controller

The I/O Controller provides control for all I/O spaces. These I/O spaces are memory mapped and are intended for use by items such as LAN controllers, SCSI controllers, I/O signal conditioners, A/D, and D/A chips. Such peripheral chips typically have address inputs, data I/O, chip select, read output enable, and if writable, a write enable strobe.

Chip Selects

The I/O Controller contains up to 8 separate memory spaces, each having its own I/O Controller Chip Select ($\overline{\text{IOCS}}$) output pin. Each $\overline{\text{IOCS}}$ space occupies from 8K to 256MB of address space of which 64MB is externally addressable (due to the 26 address lines), and the address space that each $\overline{\text{IOCS}}$ decodes is programmable. The I/O Controller will use the programmed information in the MSB and LSB Base Address Registers along with the size (8K to 256MB) of the given area as programmed in the MSB and LSB Page Mask Registers.

This information is used to compare the address asserted by the CPU-BIU or DMA Controller to determine if that particular $\overline{\text{IOCS}}$ area is being accessed for the current read or write. Each area supports single datum reads and writes. Burst reads and writes are not supported in the $\overline{\text{IOCS}}$ area. The port size of the data path (8, 16, or 32-bit) of each area is also programmable with each area's Control Register.

The $\overline{\text{IOCS}}$ signal can be used in pairs for PCMCIA. The pairs are $\overline{\text{IOCS}}(3:2)$ for one interleaved area and for the others, $\overline{\text{MemCS}}(5:4)$ and $\overline{\text{MemCS}}(7:6)$. When in the PCMCIA mode, both chip selects within a pair must be programmed to the same values.

Signal Control Interface

The I/O Controller provides read enables and write enables that are suitable for direct chip connection. I-Type read and write enables can in general also be attached to FCT543 type transceivers. M-Type write line and data strobes can in general also be attached to FCT245 type transceivers.

Wait State Generator

The Wait-State Generator (WSG) controls the speed of the I/O accesses to and from the Bus Interface Unit Controller. This includes the time from the start of an I/O transaction until the first datum is sent or received. The WSG also is programmed to generate the internal RdCEN and AckN signals for CPU read and write requests.

The internal Acknowledge signal, "AckN" (as described in Chapter 7) is the same as the external signal pin that the R3051 RISController family uses. On single word reads and on both single word and burst writes, AckN is automatically placed at the end of the transaction by the WSG. Since burst I/O reads and writes are not supported, the Control Register 'Burst Ack' field is not needed.

Since burst reads and writes to I/O are not supported, the software is required to access I/O spaces with single datum loads which are in general non-cached and the same data size as the port width.

The signal called $\overline{\text{SysWait}}$ can be used to override the programmed settings of the Wait State Generator. The actual action the WSG performs when $\overline{\text{SysWait}}$ is asserted will depend on when it is asserted relative to the transaction. $\overline{\text{SysWait}}$ has a pipeline delay, such that it must be asserted two clocks before the desired effect is noticeable. By asserting it immediately after a datum is received or transmitted, the next datum can be delayed. However, use for this purpose is generally not recommended since the WSG has the same functionality. $\overline{\text{SysWait}}$ is useful for accessing off-card "Ack"-type peripherals.

Register Option Programmability

The Memory and I/O Controller contains 8 sets of registers, one set for each chip select. These registers allow the I/O Controller to be configured for different types and speeds of peripherals. Thus almost any system speed/cost/manufacturing trade-off can be accommodated.

Register Descriptions

The Memory and I/O Controller Registers are divided into 8 sets of registers, one set for each chip select memory area. Physical addresses and register descriptions are shown in Table 9.1.

Phys. Address	Register Description
0xFFFF_E200	Memory and I/O LSB Base Address Register for Bank 0
0xFFFF_E204	Memory and I/O MSB Base Address Register for Bank 0
0xFFFF_E208	Memory and I/O LSB Bank Mask Register for Bank 0
0xFFFF_E20C	Memory and I/O MSB Bank Mask Register for Bank 0
0xFFFF_E210	Memory and I/O Control Register for Bank 0
0xFFFF_E214	Memory and I/O LSB Wait State Generator Register for Bank 0
0xFFFF_E218	Memory and I/O MSB Wait State Generator Register for Bank 0
0xFFFF_E220	Memory and I/O LSB Base Address Register for Bank 1
0xFFFF_E224	Memory and I/O MSB Base Address Register for Bank 1
0xFFFF_E228	Memory and I/O MSB LSB Bank Mask Register for Bank 1
0xFFFF_E22C	Memory and I/O MSB MSB Bank Mask Register for Bank 1
0xFFFF_E230	Memory and I/O MSB Control Register for Bank 1
0xFFFF_E234	Memory and I/O MSB LSB Wait State Generator Register for Bank 1
0xFFFF_E238	Memory and I/O MSB MSB Wait State Generator Register for Bank 1
0xFFFF_E240	Memory and I/O MSB LSB Base Address Register for Bank 2
0xFFFF_E244	Memory and I/O MSB MSB Base Address Register for Bank 2
0xFFFF_E248	Memory and I/O MSB LSB Bank Mask Register for Bank 2
0xFFFF_E24C	Memory and I/O MSB MSB Bank Mask Register for Bank 2
0xFFFF_E250	Memory and I/O MSB Control Register for Bank 2
0xFFFF_E254	Memory and I/O MSB LSB Wait State Generator Register for Bank 2
0xFFFF_E258	Memory and I/O MSB MSB Wait State Generator Register for Bank 2
0xFFFF_E260	Memory and I/O MSB LSB Base Address Register for Bank 3
0xFFFF_E264	Memory and I/O MSB MSB Base Address Register for Bank 3
0xFFFF_E268	Memory and I/O MSB LSB Bank Mask Register for Bank 3
0xFFFF_E26C	Memory and I/O MSB MSB Bank Mask Register for Bank 3
0xFFFF_E270	Memory and I/O MSB Control Register for Bank 3
0xFFFF_E274	Memory and I/O MSB LSB Wait State Generator Register for Bank 3
0xFFFF_E278	Memory and I/O MSB MSB Wait State Generator Register for Bank 3
0xFFFF_E280	Memory and I/O MSB LSB Base Address Register for Bank 4
0xFFFF_E284	Memory and I/O MSB MSB Base Address Register for Bank 4
0xFFFF_E288	Memory and I/O MSB LSB Bank Mask Register for Bank 4
0xFFFF_E28C	Memory and I/O MSB MSB Bank Mask Register for Bank 4
0xFFFF_E290	Memory and I/O MSB Control Register for Bank 4
0xFFFF_E294	Memory and I/O MSB LSB Wait State Generator Register for Bank 4
0xFFFF_E298	Memory and I/O MSB MSB Wait State Generator Register for Bank 4
0xFFFF_E2A0	Memory and I/O MSB LSB Base Address Register for Bank 5
0xFFFF_E2A4	Memory and I/O MSB MSB Base Address Register for Bank 5
0xFFFF_E2A8	Memory and I/O MSB LSB Bank Mask Register for Bank 5
0xFFFF_E2AC	Memory and I/O MSB MSB Bank Mask Register for Bank 5
0xFFFF_E2B0	Memory and I/O MSB Control Register for Bank 5
0xFFFF_E2B4	Memory and I/O MSB LSB Wait State Generator Register for Bank 5
0xFFFF_E2B8	Memory and I/O MSB MSB Wait State Generator Register for Bank 5
0xFFFF_E2C0	Memory and I/O MSB LSB Base Address Register for Bank 6
0xFFFF_E2C4	Memory and I/O MSB MSB Base Address Register for Bank 6
0xFFFF_E2C8	Memory and I/O MSB LSB Bank Mask Register for Bank 6
0xFFFF_E2CC	Memory and I/O MSB MSB Bank Mask Register for Bank 6
0xFFFF_E2D0	Memory and I/O MSB Control Register for Bank 6
0xFFFF_E2D4	Memory and I/O MSB LSB Wait State Generator Register for Bank 6
0xFFFF_E2D8	Memory and I/O MSB MSB Wait State Generator Register for Bank 6
0xFFFF_E2E0	Memory and I/O MSB LSB Base Address Register for Bank 7
0xFFFF_E2E4	Memory and I/O MSB MSB Base Address Register for Bank 7
0xFFFF_E2E8	Memory and I/O MSB LSB Bank Mask Register for Bank 7
0xFFFF_E2EC	Memory and I/O MSB MSB Bank Mask Register for Bank 7
0xFFFF_E2F0	Memory and I/O MSB Control Register for Bank 7
0xFFFF_E2F4	Memory and I/O MSB LSB Wait State Generator Register for Bank 7
0xFFFF_E2F8	Memory and I/O MSB MSB Wait State Generator Register for Bank 7

Note: Big Endian software must offset these addresses by b'10 (0x2).

Table 9.1 Memory and I/O Controller Register Addresses and Descriptions

Memory and I/O Control Register7.. (‘MemIoCntrlReg(7..0)’)

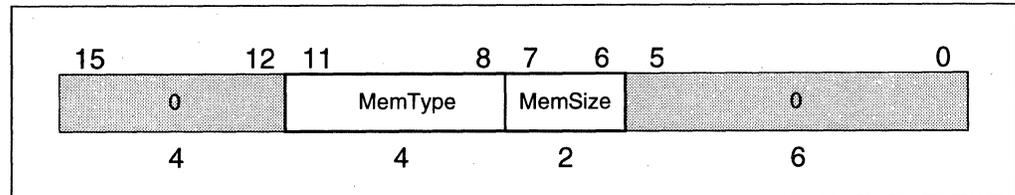


Figure 9.2 Memory and I/O Control Register (‘MemIoCntrlReg(7..0)’)

As shown in Figure 9.2, each MemCntrlReg contains a Type Field that has options specific to I/O types of signals.

Memory Type (‘MemType’) Field

The Type field determines the type of timing the Bus Interface will use. Table 9.2 lists the possible Type field values and their actions.

Value	Action
‘1011’	PCMCIA-Memory-Style
‘1010’	PCMCIA-I/O-Style
‘1001’	M-Type I/O
‘1000’	I-Type I/O
‘0010’	Memory-Type for FCT260 (default)
‘0001’	Memory-Type for FCT245 (non-interleaved)
‘0000’	Memory-Type (for FCT543)
All other values	Reserved
Note:	
PCMCIA-Style supports a PCMCIA host mode subset that is likely to be used with PCMCIA peripherals. PCMCIA-Memory and -IO Styles are intended for dynamic swapping by the software onto the same pair of chip selects. Typically, the Memory-Style is left on, and the I/O-Style is swapped in whenever it is needed, then swapped back to Memory-Style.	

Table 9.2 MemCntrlReg Memory-Type Field (‘MemType’)

Portsize Width (‘MemSize’) Field

The PortSize field, shown in Table 9.3, determines the width of the memory or I/O port. The value is inverted relative to the reset initialization vector value.

Value	Action
‘11’	64-bit (32-bit 2-way interleaved) accesses (Valid for Memory Type only)
‘10’	16-bit accesses
‘01’	8-bit accesses
‘00’	32-bit accesses

Table 9.3 PortSize (‘MemSize’) Encoding

I-Type I/O Type:

The I-Type (Intel Type) puts the bus interface into a mode such that the I/O signals $\overline{\text{IoCS}}()$, $\overline{\text{IoRd}}$, and $\overline{\text{IoWr}}$ support I-type devices. Such devices have an address decoded chip select, $\overline{\text{IoCS}}()$, and separate read ($\overline{\text{IoRd}}$) and write ($\overline{\text{IoWr}}$) data strobes. Older I-Type devices may also have an active high Reset input; an artifact that may require an external inverter.

M-Type I/O Type:

The M-Type (Motorola Type) puts the bus interface into a mode such that the I/O signals $\overline{\text{IoRd}}$ and $\overline{\text{IoWr}}$ are used as $\overline{\text{IoDStrobe}}$ and $\overline{\text{IoRdHWr}}$, respectively. The device is assumed to have an address decoded chip select, $\overline{\text{IoCS}}()$ while using the $\overline{\text{IoRdHWr}}$ line as a status line indicating a read or write, and using $\overline{\text{IoDStrobe}}$ as a data strobe.

PCMCIA-I/O Style:

The PCMCIA-Style puts the bus interface into a mode such that the I/O signals $\overline{\text{IoCS}}(\text{odd,even})$, $\overline{\text{IoRd}}$, and $\overline{\text{IoWr}}$ support 16-bit slave PCMCIA devices. The odd and even chip select pair is used to indicate whether one or both of the byte lanes are valid. $\overline{\text{IoRd}}$ and $\overline{\text{IoWr}}$ are used as separate read and write data strobes.

I/O Controller Timing Diagrams

This section illustrates a number of timing diagrams applicable to R36100 I/O transactions. The values for the AC parameters are contained in the separate document, "R36100 RISController Data Sheet."

I/O Datum Size

All I/O accesses must be single datum accesses. For example, a 32-bit port must not use a cached burst access; a 16-bit port must use an uncached store or load halfword or byte operation (such as one of these: sh, lhu, sb, lbu). An 8-bit port must only use uncached byte operations (such as one of these: sb, lbu).

Read Transactions

The bus interface timing for read transactions is described in the following section. The internal bus interface to CPU core for read loads is described in Chapter 7.

Basic I-Type I/O Read with 0 Wait-States

Figure 9.3 illustrates a basic I-Type I/O Controller read transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysRd}}$ asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that $\overline{\text{IoCS()}}$, $\overline{\text{IoRd}}$, and $\overline{\text{IoWr}}$ will be in their de-asserted states. And assuming there are no internally programmed StartRepeat wait-states, on the next clock cycle, $\overline{\text{SysBurstFrame}}$ de-asserts and $\overline{\text{IoCS()}}$ asserts.

On the third cycle, $\overline{\text{IoRd}}$ asserts. Then assuming there are no internally programmed RdStart2Datum wait-states, $\overline{\text{SysDataRdy}}$ asserts to indicate that the data from the I/O device is being sampled into the RISController. On the 4th clock cycle, $\overline{\text{IoRd}}$ de-asserts, indicating that the read data from the I/O device has just been latched into the RISController. On the next clock--the final clock of the transaction-- $\overline{\text{IoCS()}}$ de-asserts, and the next transaction may begin.

During an I-Type transaction, there is 1 clock of address setup time before $\overline{\text{IoCS()}}$ asserts. All signals are setup before the read strobe, $\overline{\text{IoRd}}$, asserts. After data has been sampled by the CPU, $\overline{\text{IoRd}}$ de-asserts, with all other signals having hold time. On the next clock, $\overline{\text{IoCS()}}$ de-asserts.

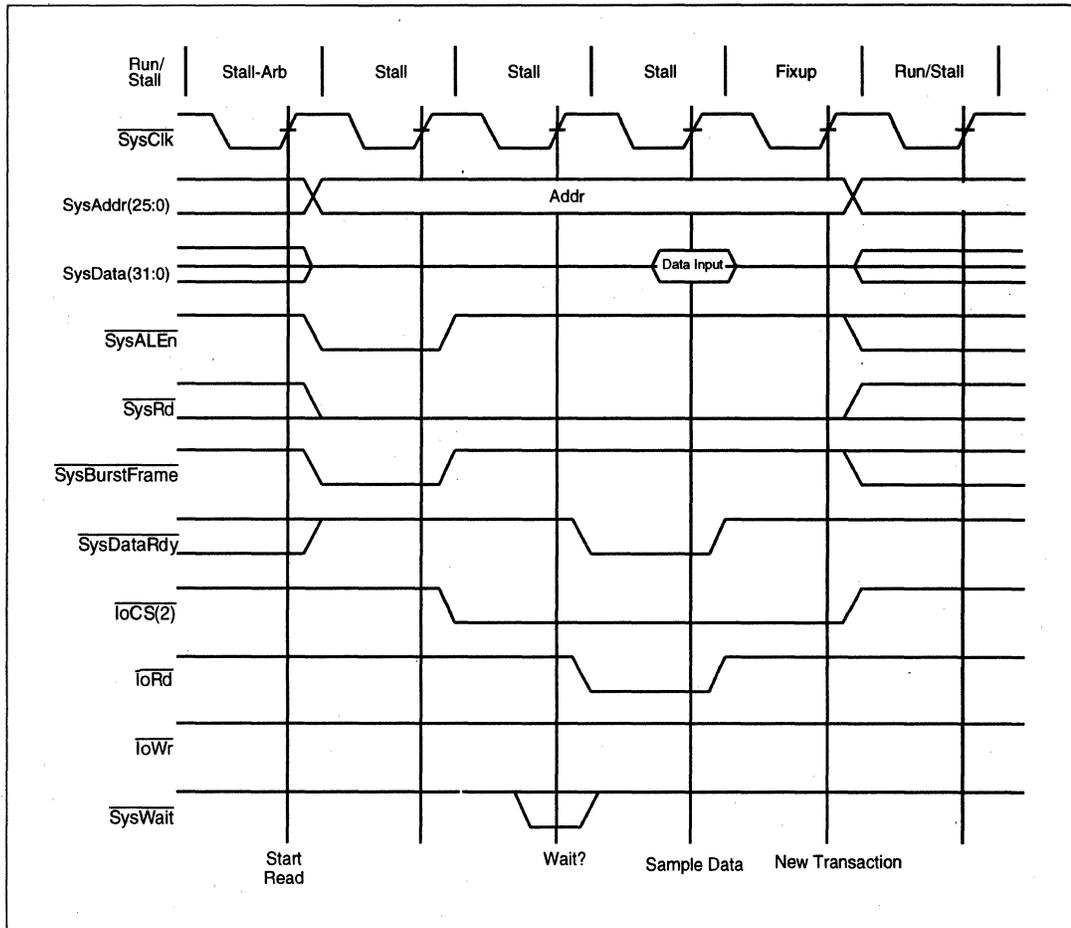


Figure 9.3 I-Type I/O Read with 0 Wait-States

Basic M-Type I/O Read with 0 Wait-States

Figure 9.4 illustrates a basic M-Type I/O Controller read transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysRd}}$ asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that $\overline{\text{IoCS}(0)}$, $\overline{\text{IoDStrobe}}$ (also known as $\overline{\text{IoRd}}$), and $\overline{\text{IoRdHWr}}$ (also known as $\overline{\text{IoWr}}$) will be in their de-asserted states. Then assuming there are no internally programmed StartRepeat wait-states, on the next clock cycle, $\overline{\text{SysBurstFrame}}$ de-asserts; and $\overline{\text{IoCS}(0)}$ asserts. If this is a write transaction, $\overline{\text{IoRdHWr}}$ ($\overline{\text{IoWr}}$) will assert.

On the third cycle, $\overline{\text{IoDStrobe}}$ ($\overline{\text{IoRd}}$) asserts. And assuming there are no internally programmed RdStart2Datum wait-states, $\overline{\text{SysDataRdy}}$ asserts, to indicate that the data from the I/O device is being sampled into the RISController. On the 4th clock cycle, $\overline{\text{IoDStrobe}}$ ($\overline{\text{IoRd}}$) de-asserts, indicating that the read data from the I/O device has just been latched into the RISController. On the next clock--the final clock of the transaction-- $\overline{\text{IoCS}(0)}$ de-asserts and, at the same time, the next transaction may begin.

During an M-Type transaction, there is 1 clock of address setup time before $\overline{\text{IoCS}(0)}$ asserts. All signals are setup before the data strobe, $\overline{\text{IoDStrobe}}$, asserts. After data has been sampled by the CPU, $\overline{\text{IoDStrobe}}$ de-asserts, with all other signals having hold time. On the next clock, $\overline{\text{IoCS}(0)}$ de-asserts.

Note that because $\overline{\text{IoRdHWr}}$ asserts on the same clock edge as $\overline{\text{IoCS}(0)}$, systems that require setup time from $\overline{\text{IoRdHWr}}$ to $\overline{\text{IoCs}(0)}$ can substitute $\overline{\text{SysWr}}$ for $\overline{\text{IoRdHWr}}$. Many systems can also substitute $\overline{\text{IoCS}(0)}$ for $\overline{\text{IoDStrobe}}$ which, if the peripheral timing allows, may require fewer wait-states.

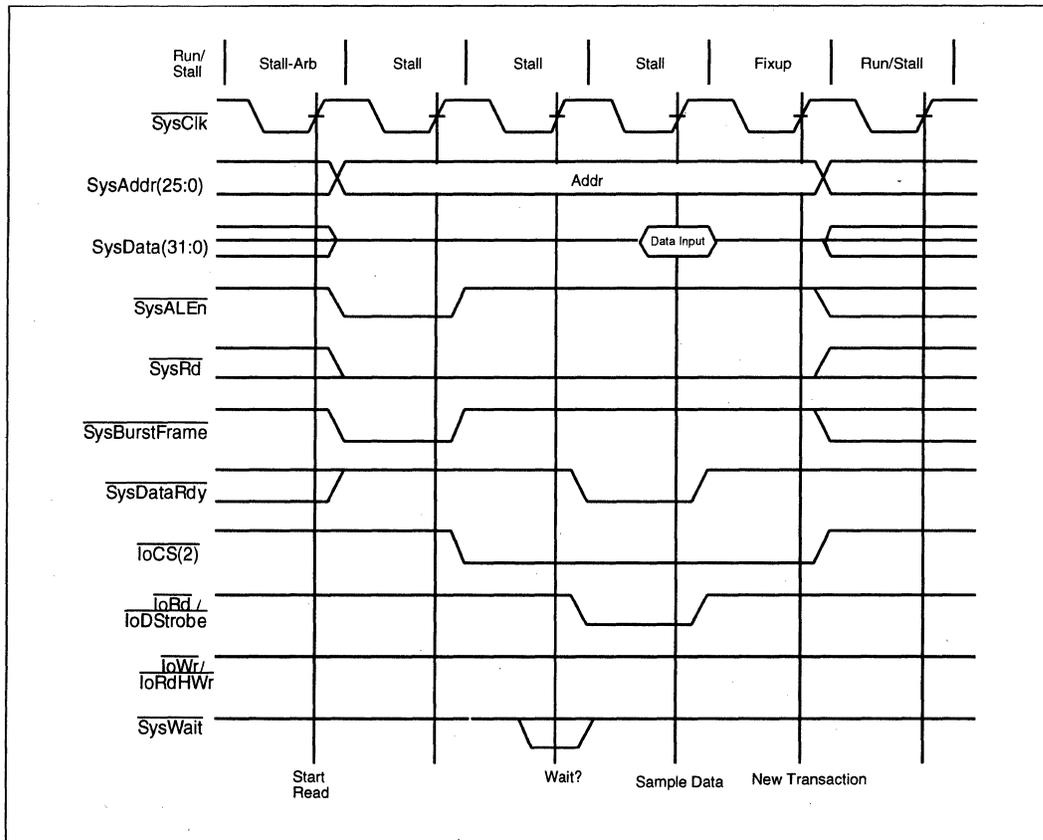


Figure 9.4 M-Type I/O Read with 0 Wait-States

Basic 16-bit PCMCIA-style I/O Read with 0 Wait-States

Figure 9.5 illustrates a basic 16-bit PCMCIA-Style I/O Controller read transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysRd}}$ asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that $\overline{\text{IoCS0}}$, $\overline{\text{IoRd}}$, and $\overline{\text{IoWr}}$ will be in their de-asserted states. Then assuming that there are no internally programmed StartRepeat wait-states, on the next clock cycle, $\overline{\text{SysBurstFrame}}$ de-asserts and the $\overline{\text{IoCS0}}$ pair asserts. Note that on PCMCIA transactions, the $\overline{\text{IoCS0}}$ pair is asserted according to which bytes are enabled and valid. Thus if the even byte is valid, then the even $\overline{\text{IoCS0}}$ will assert. If the odd byte is valid, then the odd $\overline{\text{IoCS0}}$ will assert. If both bytes are valid, then both $\overline{\text{IoCS0}}$ signals in the pair assert.

On the third cycle $\overline{\text{IoRd}}$ asserts. And assuming there are no internally programmed RdStart2Datum wait-states, $\overline{\text{SysDataRdy}}$ asserts to indicate that the data from the I/O device is being sampled into the RISController. On the 4th clock cycle, $\overline{\text{IoRd}}$ de-asserts, indicating that the read data from the I/O device has just been latched into the RISController. On the next clock--the final clock of the transaction--the $\overline{\text{IoCS0}}$ pair de-asserts and, at the same time, the next transaction may begin.

During an I-Type transaction, there is 1 clock of address setup time before the $\overline{\text{IoCS0}}$ pair asserts. All signals are setup before the read strobe, $\overline{\text{IoRd}}$, asserts. After data has been sampled by the CPU, $\overline{\text{IoRd}}$ de-asserts, with all other signals having hold time. On the next clock the $\overline{\text{IoCS0}}$ pair de-asserts.

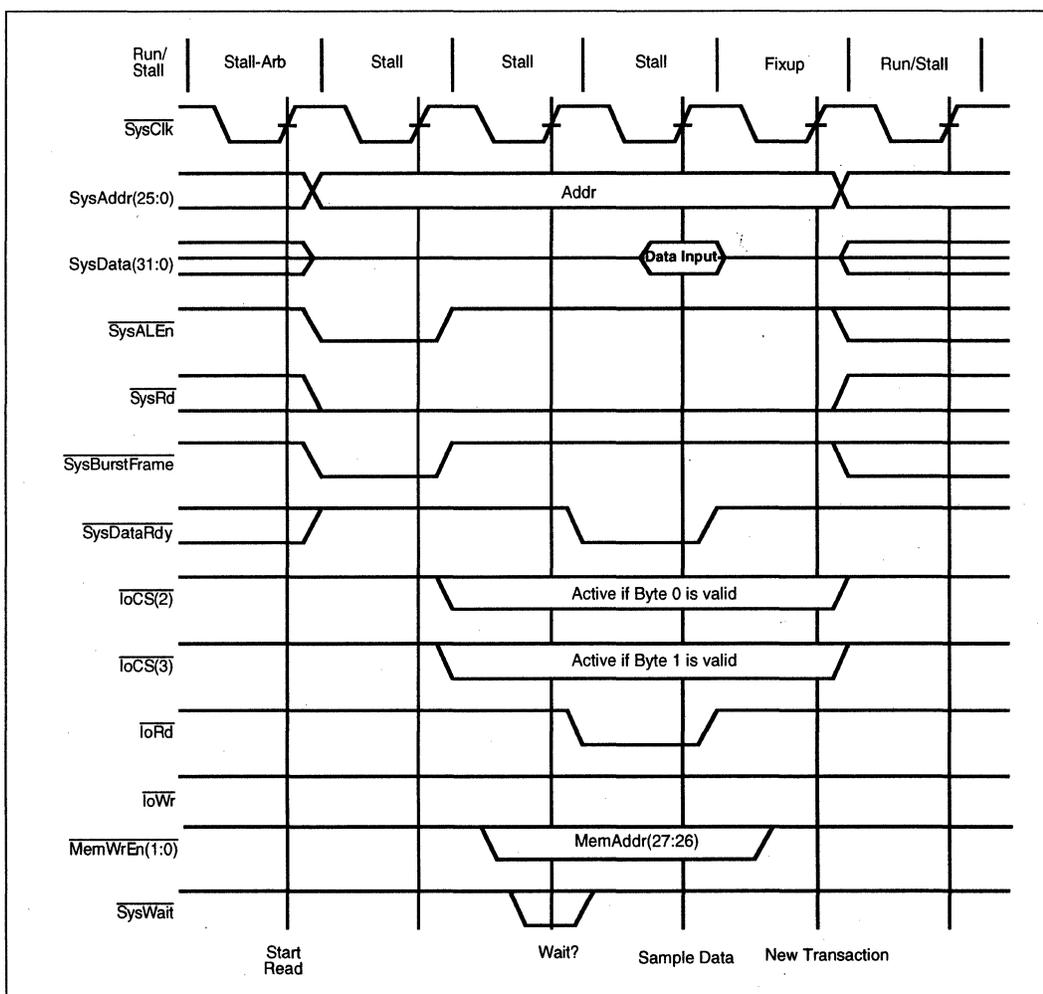


Figure 9.5 PCMCIA-Style I/O Read with 0 Wait-States

Basic I-Type I/O Write with 0 Wait-States

Figure 9.6 illustrates a basic I-Type I/O Controller write transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysWr}}$ asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that $\overline{\text{IoCS}(0)}$, $\overline{\text{IoRd}}$, and $\overline{\text{IoWr}}$ will be in their de-asserted states. Assuming that there are no internally programmed StartRepeat wait-states, on the next clock cycle, $\overline{\text{SysBurstFrame}}$ de-asserts and $\overline{\text{IoCS}(0)}$ asserts.

On the third cycle, $\overline{\text{IoWr}}$ asserts. And assuming there are no internally programmed RdStart2Datum wait-states, $\overline{\text{SysDataRdy}}$ asserts to indicate that the data from the RISController is ready to be latched into the I/O device. On the 4th clock cycle, $\overline{\text{IoWr}}$ de-asserts, providing a means for the write data from the RISController to be latched into the I/O device. On the next clock--which is the final clock of the transaction-- $\overline{\text{IoCS}(0)}$ de-asserts and, at the same time, the next transaction may begin.

During an I-Type transaction, there is 1 clock of address setup time before $\overline{\text{IoCS}(0)}$ asserts. All signals are setup before the write strobe, $\overline{\text{IoWr}}$, asserts. After data has been sampled by the CPU, $\overline{\text{IoWr}}$ de-asserts, with all other signals having hold time. On the next clock, $\overline{\text{IoCS}(0)}$ de-asserts.

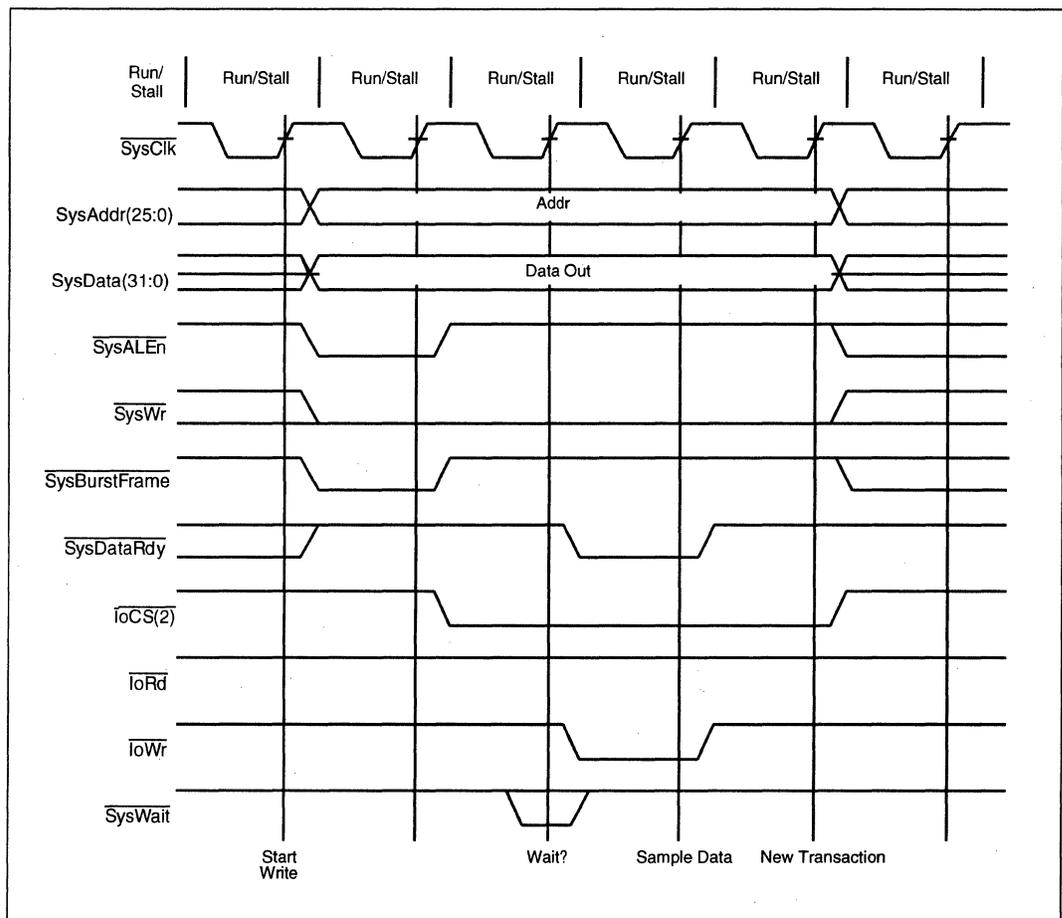


Figure 9.6 I-Type I/O Write with 0 Wait-States

Basic M-Type I/O Write with 0 Wait-States

Figure 9.7 illustrates a basic M-Type I/O Controller write transaction. Each transaction begins with both $\overline{\text{SysALEn}}$ and $\overline{\text{SysBurstFrame}}$ asserting. At this time, $\overline{\text{SysWr}}$ asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that $\overline{\text{IoCS}(0)}$, $\overline{\text{IoDStrobe}}$ (also known as $\overline{\text{IoRd}}$), and $\overline{\text{IoRdHWr}}$ (also known as $\overline{\text{IoWr}}$) will be in their de-asserted states. Assuming that there are no internally programmed StartRepeat wait-states, on the next clock cycle, $\overline{\text{SysBurstFrame}}$ de-asserts; and $\overline{\text{IoCS}(0)}$ asserts. $\overline{\text{IoRdHWr}}$ ($\overline{\text{IoWr}}$) also asserts here to indicate the I/O write transaction.

On the third clock cycle $\overline{\text{IoDStrobe}}$ ($\overline{\text{IoRd}}$) asserts. And assuming that there are no internally programmed RdStart2Datum wait-states, $\overline{\text{SysDataRdy}}$ asserts to indicate that the data from the I/O device is being sampled into the RISController. On the 4th clock cycle, $\overline{\text{IoDStrobe}}$ ($\overline{\text{IoRd}}$) de-asserts indicating that the read data from the I/O device has just been latched into the RISController. On the next clock--the final clock of the transaction-- $\overline{\text{IoCS}(0)}$ and $\overline{\text{IoRdHWr}}$ de-assert and, at the same time, the next transaction may begin.

During an M-Type transaction, there is 1 clock of address setup time before $\overline{\text{IoCS}(0)}$ asserts. All signals are setup before the data strobe, $\overline{\text{IoDStrobe}}$, asserts. After data has been sampled by the CPU, $\overline{\text{IoDStrobe}}$ de-asserts with all other signals having hold time. On the next clock, $\overline{\text{IoCS}(0)}$ de-asserts.

Note that because $\overline{\text{IoRdHWr}}$ asserts on the same clock edge as $\overline{\text{IoCS}(0)}$, systems that require setup time from $\overline{\text{IoRdHWr}}$ to $\overline{\text{IoCs}(0)}$ can substitute $\overline{\text{SysWr}}$ for $\overline{\text{IoRdHWr}}$. Many systems can also substitute $\overline{\text{IoCS}(0)}$ for $\overline{\text{IoDStrobe}}$, which if the peripheral timing allows, may require fewer wait-states.

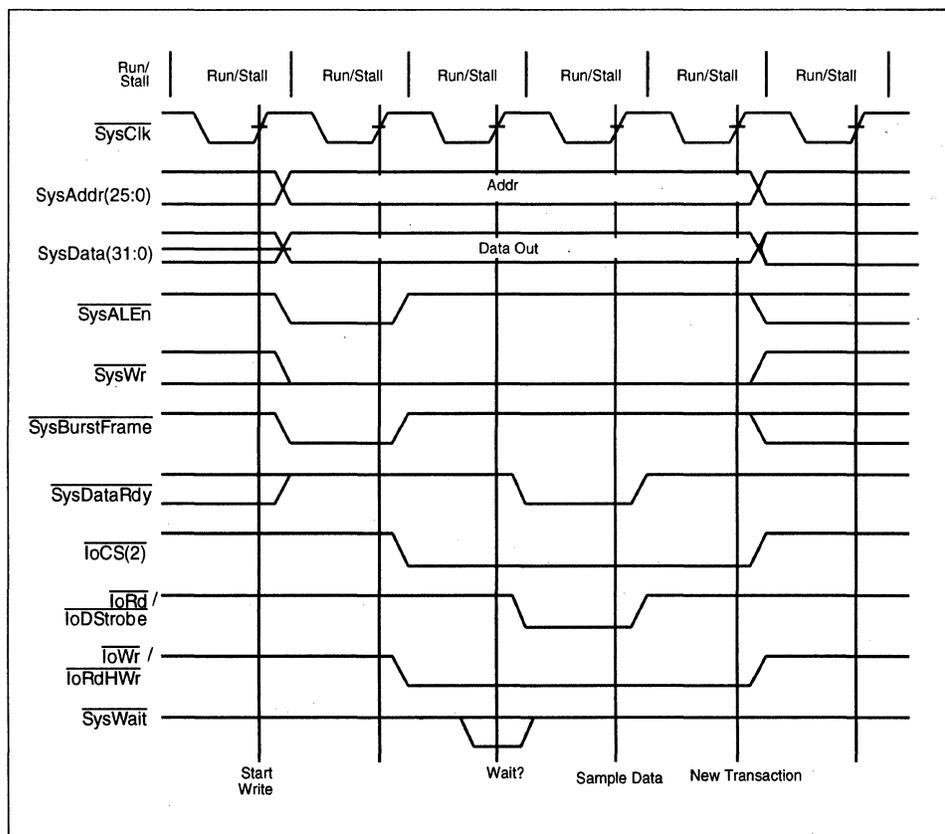


Figure 9.7 M-Type I/O Write with 0 Wait-States

Basic 16-bit PCMCIA-Style I/O Write with 0 Wait-States

Figure 9.8 illustrates a basic 16-bit PCMCIA-Style I/O Controller write transaction. Each transaction begins with both SysALEn and SysBurstFrame asserting. At this time, SysWr asserts (if it is not already in this state, as the result of a previous transaction). Now it is guaranteed that IoCS0, IoRd, and IoWr will be in their de-asserted states. And assuming there are no internally programmed StartRepeat wait-states, on the next clock cycle, SysBurstFrame de-asserts, and the IoCS0 pair asserts. Note that on PCMCIA transactions the IoCS0 pair is asserted according to which bytes are enabled and valid. Thus if the even byte is valid, then the even IoCS0 will assert. If the odd byte is valid, then the odd IoCS0 will assert. If both bytes are valid, then both IoCS0 signals in the pair assert.

On the third cycle IoWr asserts. Assuming there are no internally programmed RdStart2Datum wait-states, SysDataRdy asserts to indicate that the data from the RISController is ready to be latched into the I/O device. On the fourth clock cycle, IoWr de-asserts, providing a means for the write data from the RISController to be latched into the I/O device. On the next clock cycle--the final clock of the transaction--the IoCS0 pair de-asserts and, at the same time, the next transaction may begin.

During an I-Type transaction, there is 1 clock of address setup time before the IoCS0 pair asserts. All signals are setup before the write strobe, IoWr, asserts. After data has been sampled by the CPU, IoWr de-asserts with all other signals having hold time. On the next clock, the IoCS0 pair de-asserts.

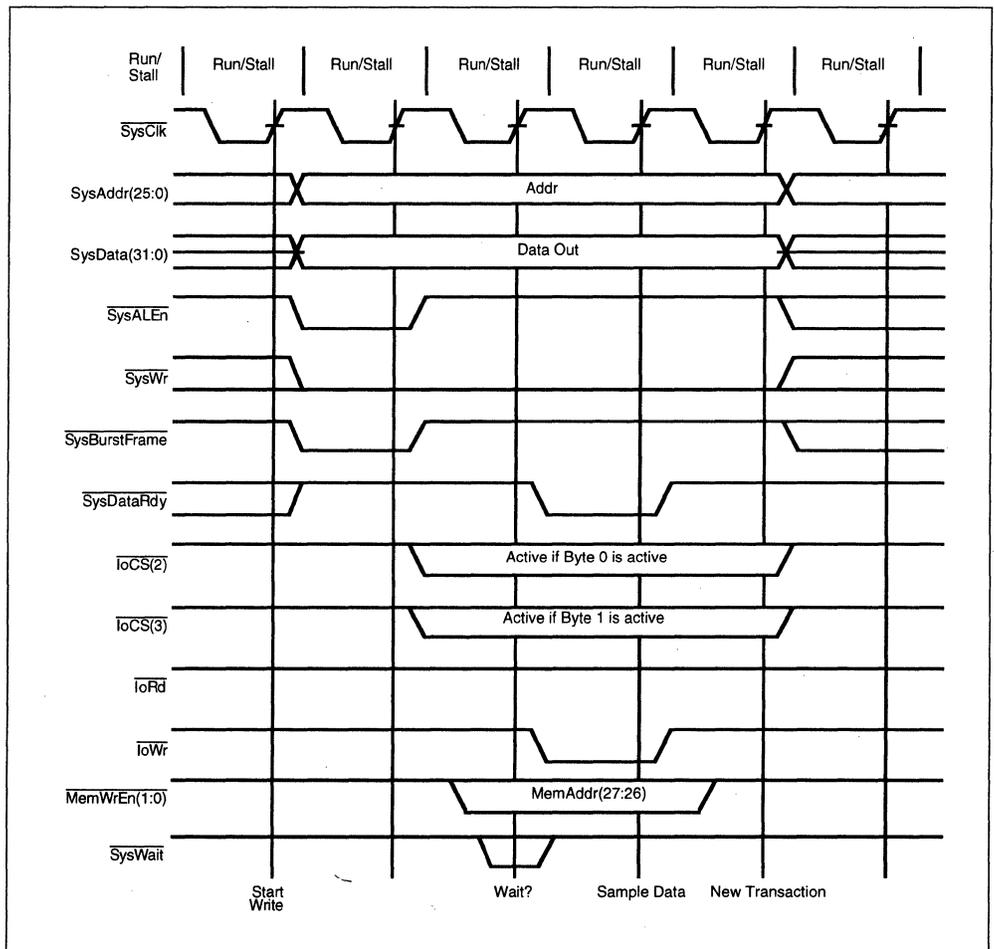


Figure 9.8 PCMCIA-Style I/O Write with 0 Wait-States

Read with Wait-State using Start Repeat Field

The left half of Figure 9.9 illustrates a basic I/O Controller read where 1 wait-state has been added by repeating the Start Cycle. Although Figure 9.9 shows an I-Type transaction, the same general timing concept applies to M-Type, and PCMCIA-style accesses as well. This special effect is programmed into the Wait-State Generator using the Start Repeat Field in the MemIoLSBWaitStateReg() Register. When the Start Cycle repeats, the $\overline{\text{IoCS}}()$ assertion is delayed. This effect is useful for very slow peripherals or peripherals that require significant address setup before the chip is selected. An example is the 600ns access time mode of the PCMCIA I/O protocol. The Start repeat Field affects both reads and writes.

Read with Wait-State using RdStart2Datum Field

The right half of Figure 9.9 illustrates a basic I/O Controller read where 1 wait-state is added using the RdStart2Datum Field of the MemIoMSBWaitStateReg() Register. Any number from 0 to 15 internal wait-states may be added using the RdStart2Datum Field. With this field, $\overline{\text{IoCS}}()$ and the read, write, or data strobe is asserted as normal, but then wait-states are added such that $\overline{\text{SysDataRdy}}$ is not asserted until the RdStart2Datum Field has finished counting. When $\overline{\text{SysDataRdy}}$ is asserted, then the Data from the external Memory Bank is sampled into the RISController.

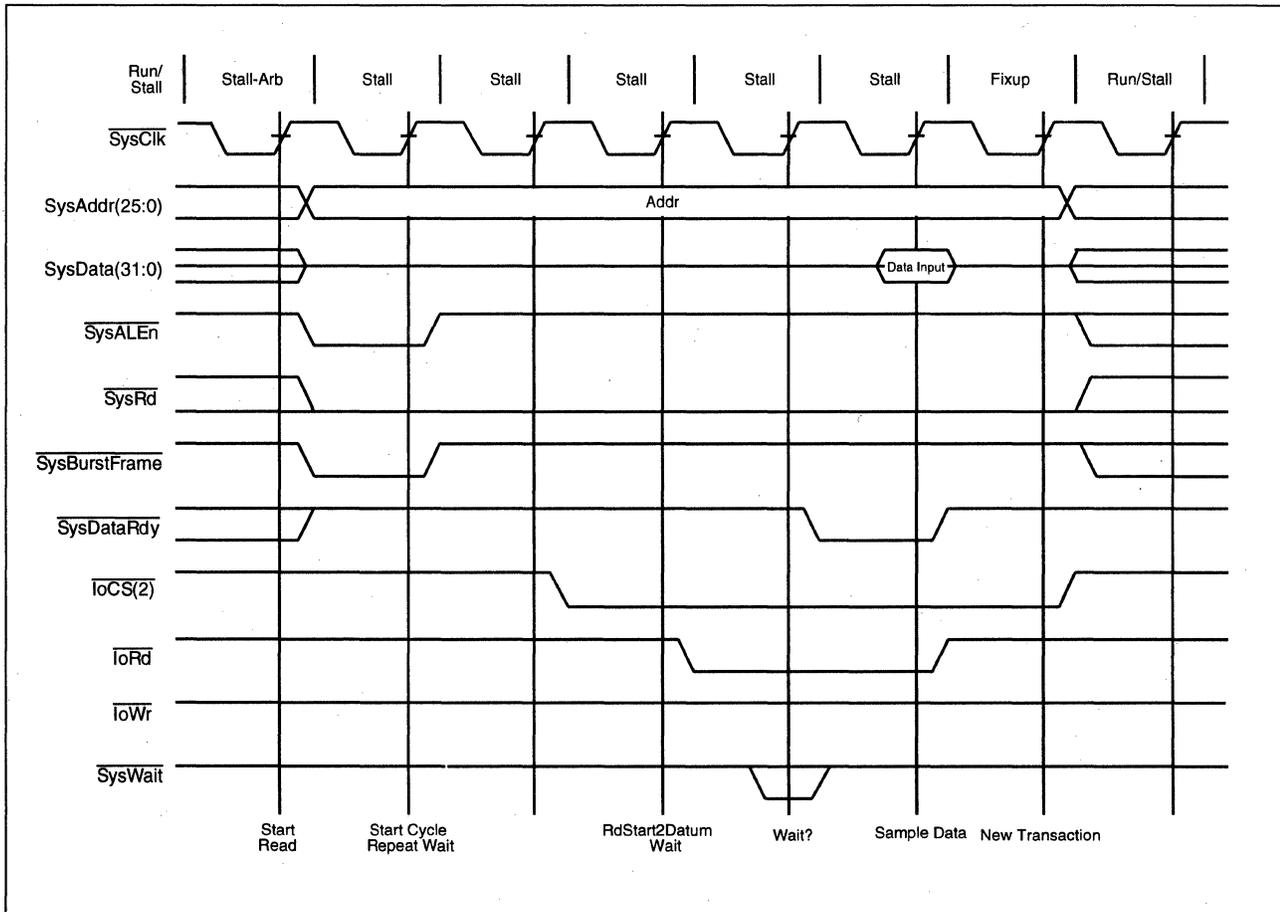


Figure 9.9 I/O Read with Internal Wait-States

Read with Wait-State using SysWait

Figure 9.10 illustrates a basic I/O Controller read where 1 wait-state is added using the external signal pin, SysWait. SysWait is not expected to be used for conventional I/O systems, since it is easier to program the Wait-State Generator to produce internal wait-states. However, SysWait can be useful for off-card I/O where there may be an indeterminate amount of time before the access can begin. Since SysWait is sampled a clock ahead of when it is used, its effect is seen two clocks later than when it is asserted. If SysWait is asserted when SysDataRdy is asserted, then an additional Data Sampling clock cycle is repeated with SysDataRdy remaining low. Thus external logic analyzers or other debug equipment may want to gate SysDataRdy with SysWait in order to decode valid Data samples.

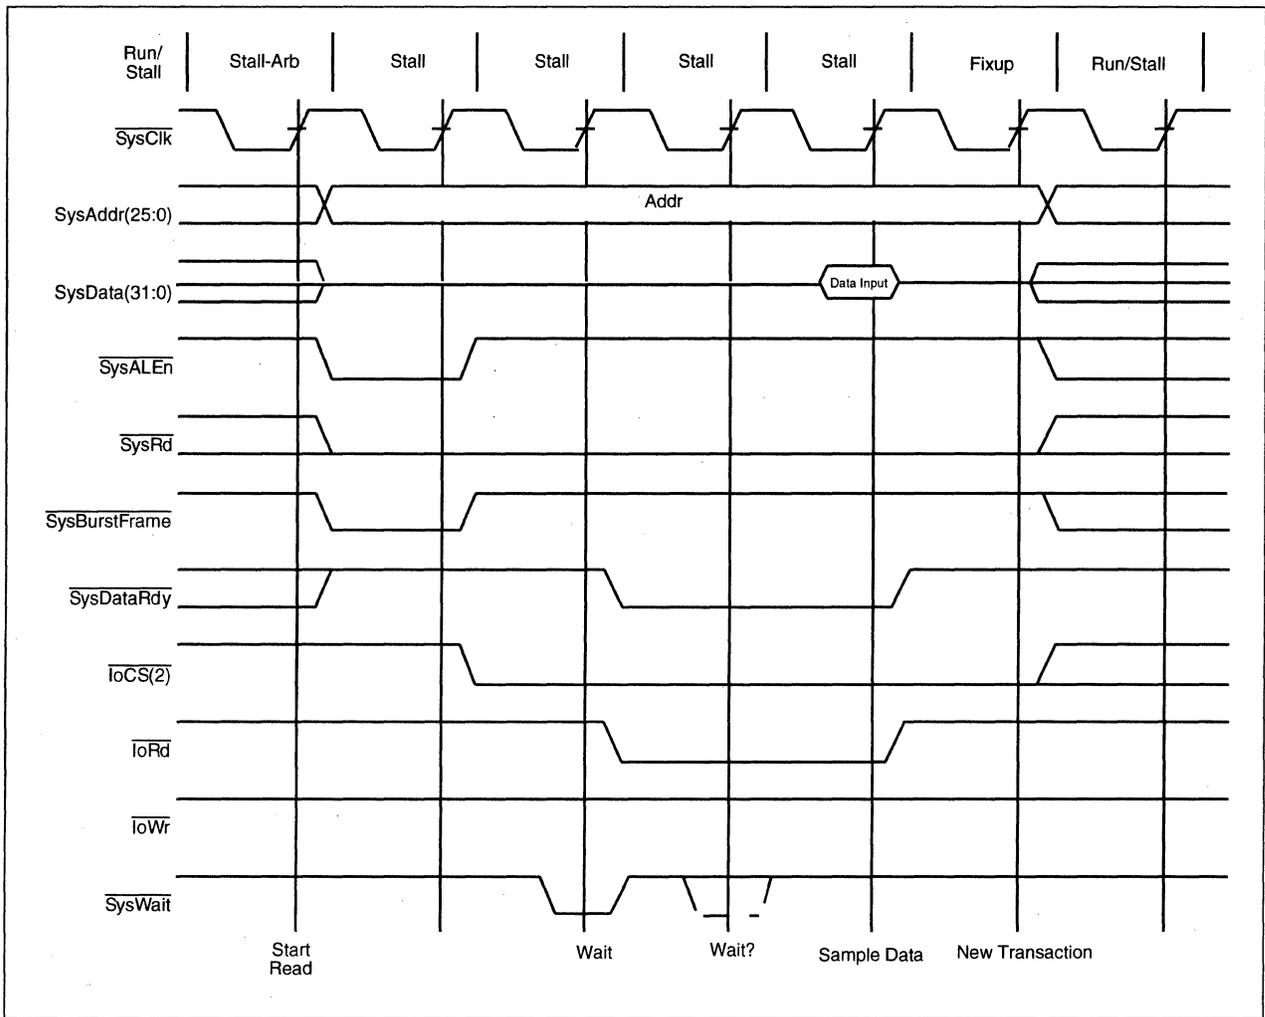


Figure 9.10 I/O Read with external SysWait Wait-State

1-Datum Write with Wait-State using StartRepeat Field

The right half of Figure 9.11 illustrates a basic I/O Controller write where 1 wait-state has been added by repeating the Start Cycle. This special effect is programmed into the Wait-State Generator using the Start Repeat Field in the MemIoLSBWaitStateReg() Register. When the Start Cycle repeats, the assertion of $\overline{\text{IoCS}}(2)$ is delayed. This effect is useful for very slow peripherals or peripherals that require significant address setup before the chip is selected. An example is the 600ns access time mode of the PCMCIA I/O protocol. The Start repeat Field affects both reads and writes.

1-Datum Write with Wait-State using WrStart2Datum Field

The left half of Figure 9.11 illustrates a basic I/O Controller write where 1 wait-state is added using the WrStart2Datum Field of the MemIoMSBWaitStateReg() Register. Any number from 0 to 15 internal wait-states may be added using the WrStart2Datum Field. With this field, the I/O Write Enable strobe, either $\overline{\text{IoWr}}$ or $\overline{\text{IoDStrobe}}$ is asserted as normal, but then wait-states are added where $\overline{\text{SysDataRdy}}$ is not asserted until the WrStart2Datum Field has finished counting. When $\overline{\text{SysDataRdy}}$ is asserted, then the Data from the external Memory Bank is sampled by external memory.

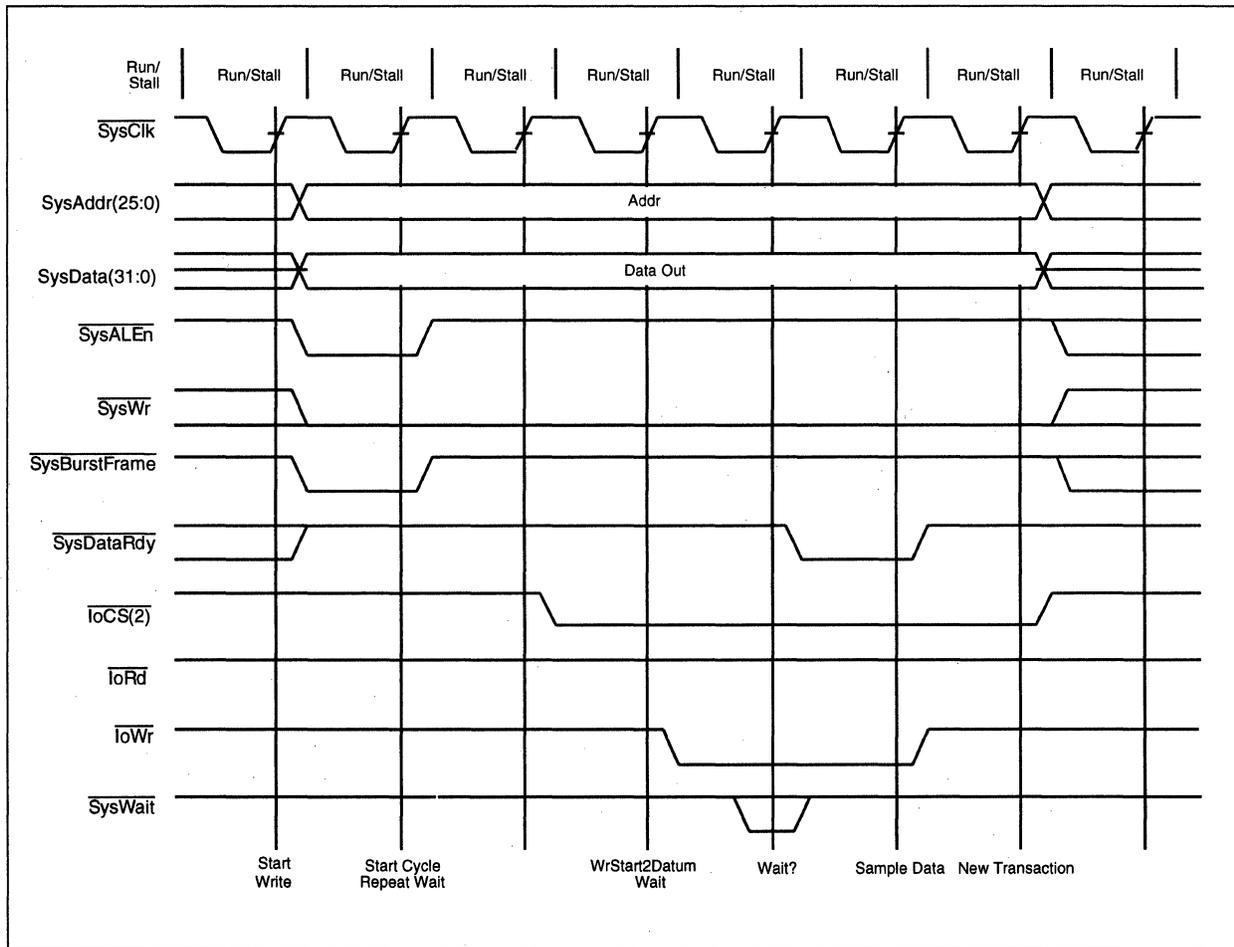


Figure 9.11 I/O Write with Internal Wait-States

1-Datum Write with Wait-State using SysWait

Figure 9.12 illustrates a basic I/O Controller write where 1 wait-state is added using the external signal pin, SysWait. SysWait is not Wait-State Generator to produce internal wait-states. However, SysWait can be useful for off-card peripherals where there may be an indeterminate amount of time before the access can begin. Since SysWait is sampled a clock ahead of when it is used, its effect is seen two clocks later than when it is asserted. If SysWait is asserted when SysDataRdy is asserted then an additional Data Sampling clock cycle is repeated with SysDataRdy remaining low. Thus external logic analyzers or other debug equipment may want to gate SysDataRdy with SysWait in order to decode valid Data samples.

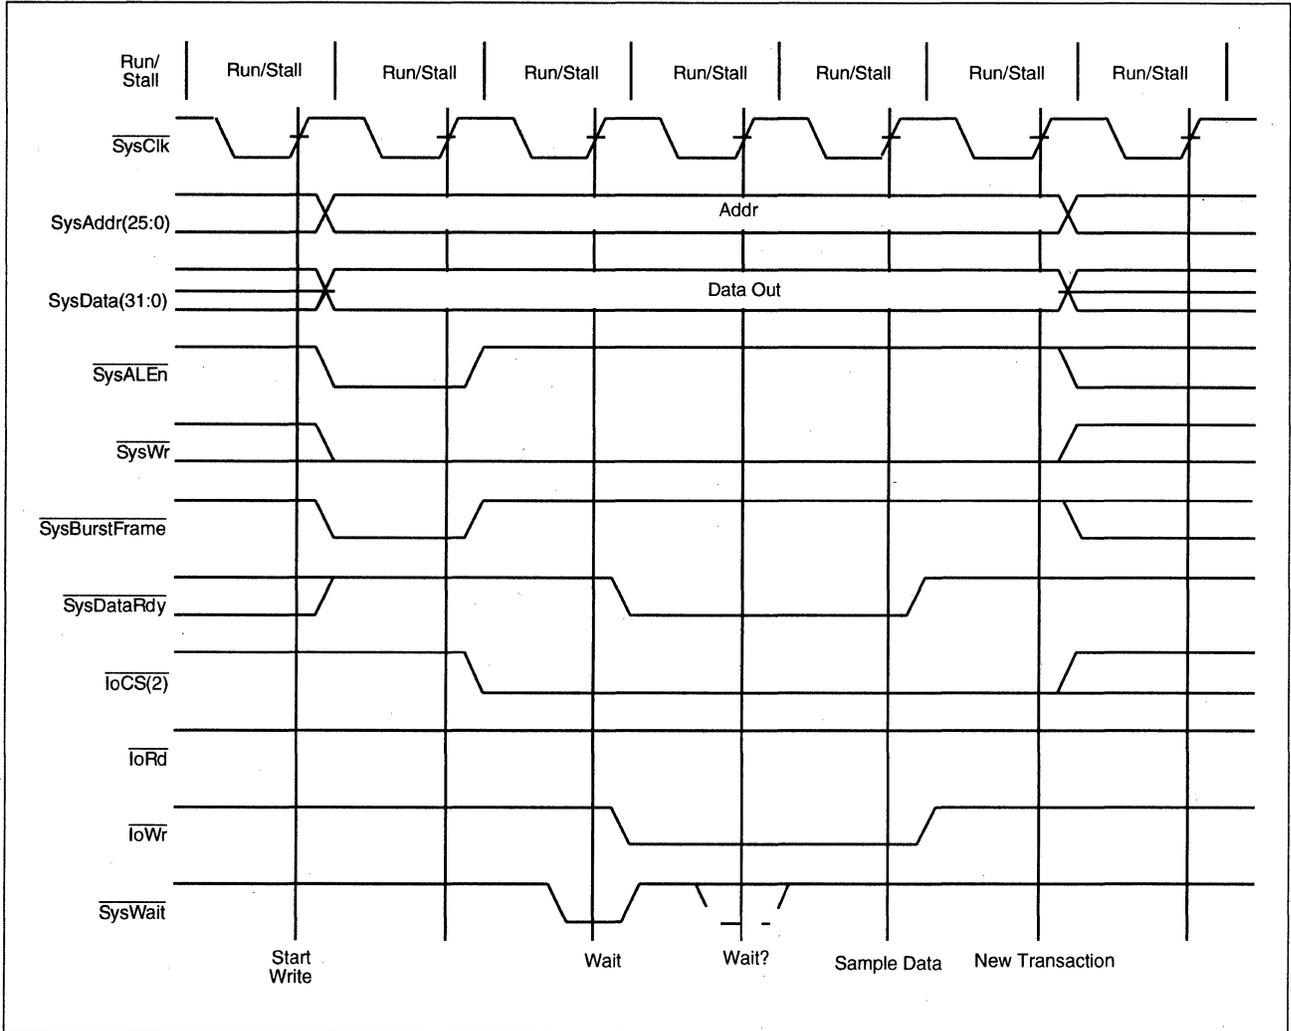


Figure 9.12 I/O Write with external SysWait Wait-State

System Examples

32-bit I/O Device Directly Connected to Bus

Figure 9.13 shows a typical 32-bit I/O device using the I-Type. And Figure 9.14 shows a typical 32-bit I/O device using the M-Type. In small systems, the I/O device can be attached directly to the SysAddr and SysData buses. If the current load is relatively large or the device turn-off time after a read is relatively long, the I/O device should be isolated with a transceiver.

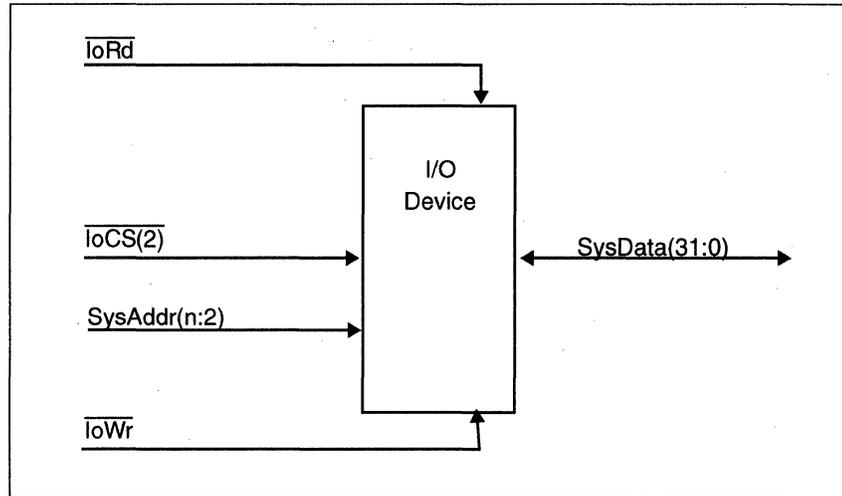


Figure 9.13 I-Type I/O System with Direct Bus Connection

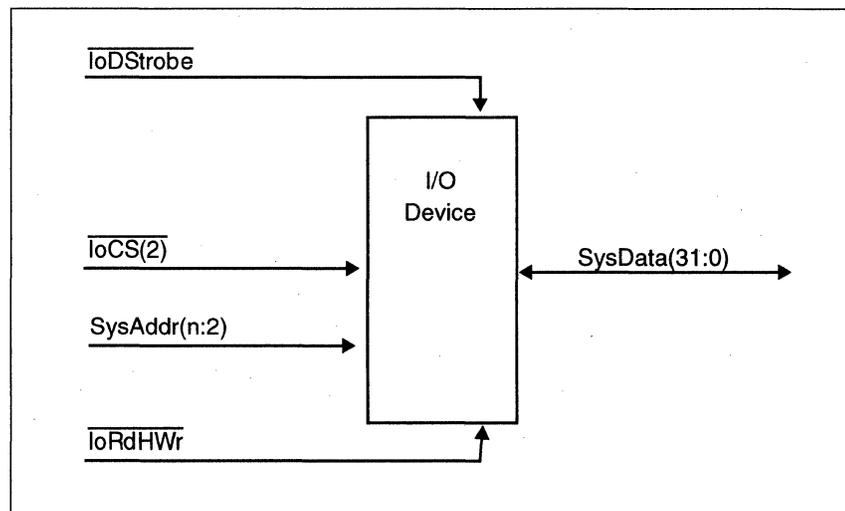


Figure 9.14 M-Type I/O System with Direct Bus Connection

I/O Reset Application

Certain types of I/O devices require both Read and Write to be asserted for generating a Reset. Type 3 can be used to do so, however, this requires that no other devices be of Type 3. If such devices must be used with Type 3 devices, then the Read and Write Reset device must externally gate its chip select with either the read or write line.

Other types of I/O devices require an active high reset (I-type). This can either be accomplished with an external inverter of ResetN or with a spare PIO line.

32-bit I/O Device using 245 Transceivers

Figure 9.15 shows a typical I-Type I/O device using FCT245 transceivers. Figure 9.16 shows a typical M-Type I/O device using FCT245 transceivers.

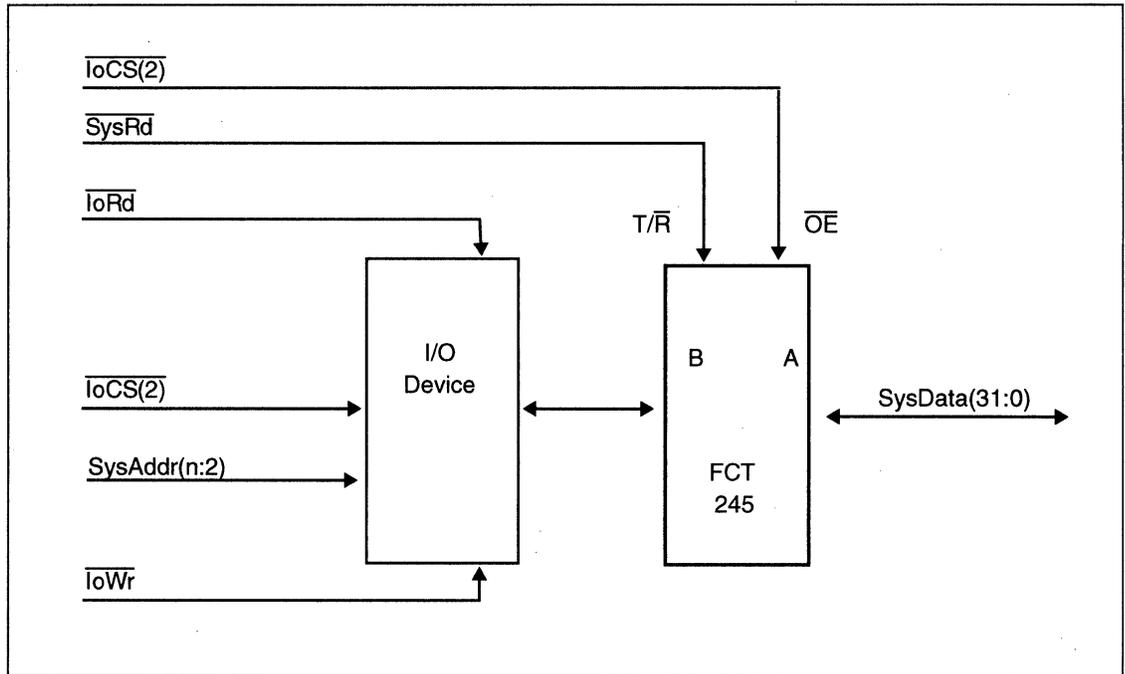


Figure 9.15 I-Type I/O System using FCT245 Transceivers

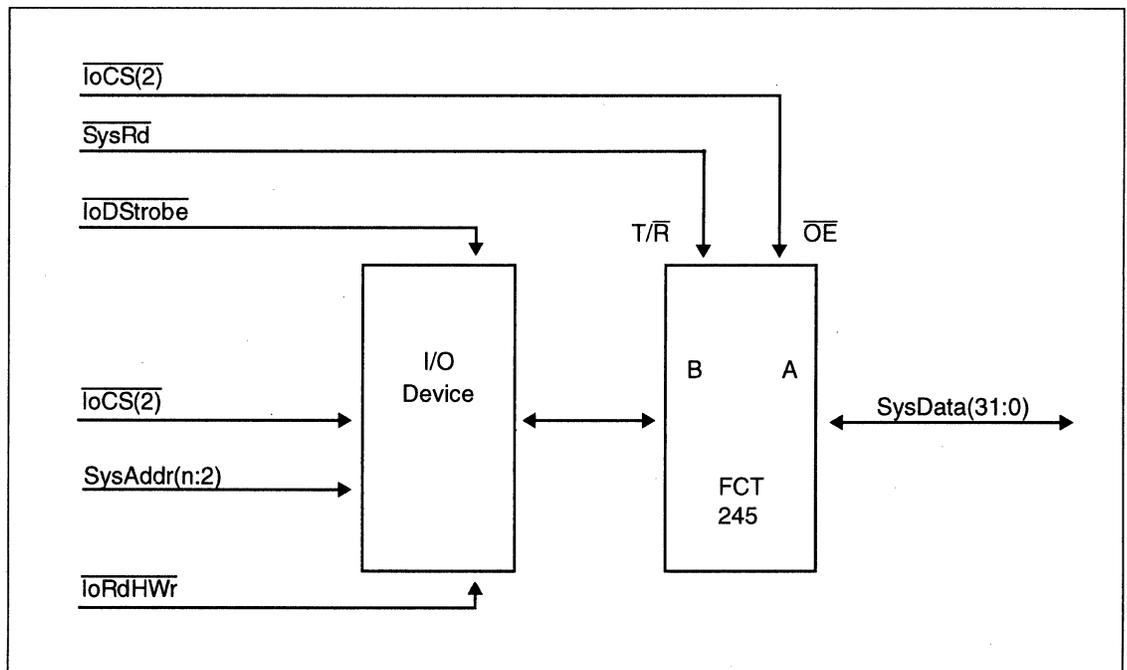


Figure 9.16 M-Type I/O System using FCT245 Transceivers

32-bit I/O Device using 543 Transceivers

Figure 9.17 shows a typical I-Type I/O device using FCT543 transceivers. This example system takes advantage of the dual output enable and chip select gating of the FCT543 where both the output enable and the chip select need to be asserted for the transceiver to drive its outputs. Figure 9.18 shows a typical M-Type I/O device using FCT543 transceivers.

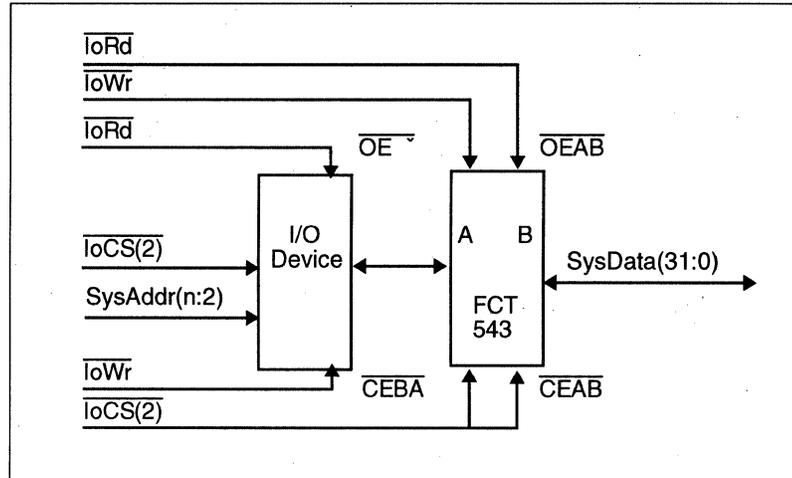


Figure 9.17 I-Type I/O System using FCT543 Transceivers

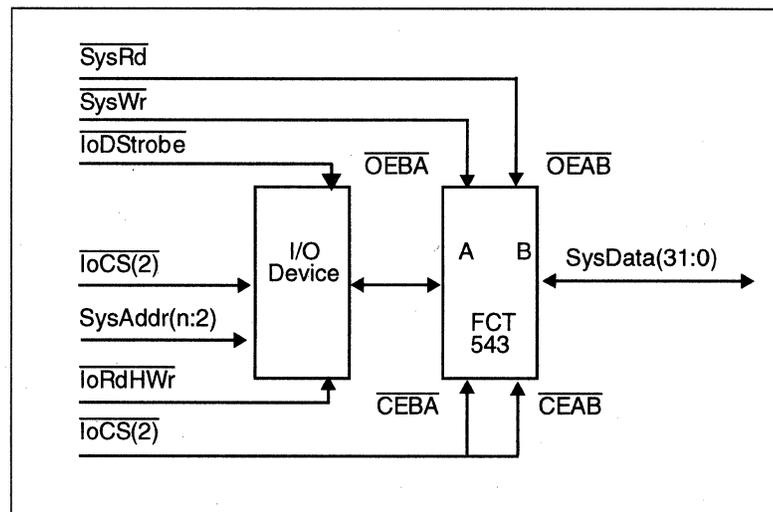


Figure 9.18 M-Type I/O System using FCT543 Transceivers

Using more than one device behind each transceiver

Multiple I/O devices can be put behind the same set of transceivers. The most obvious method is to add an external decoder to divide the chip select up into individual chip selects for each device using the MSB SysAddr line as the select. A second method is to use a spare R36100 IoCS0 pin and assign it the same address spaces as the devices behind the transceiver. Thus the common IoCS0 combines all of the address spaces of the devices behind the transceiver, such that the transceiver is turned on for any of those devices.

16-bit I/O Devices

Figure 9.19 and Figure 9.20 show typical 16-bit I/O devices. There are two choices for hooking up 16-bit I/O devices.

For the first option, refer back to the 32-bit case in Figure 9.15 on page 23 and Figure 9.16 on page 23. The 16-bit device is word-aligned (using SysAddr bits n:2) even though it is a 16-bit device. This is the traditional MIPS connection and allows the device to be accessed from either endianness. For example, if the device is connected to SysData(15:0), then little endian software accesses the registers like 0x00, 0x04, 0x08, 0x0C, ... and big endian software accesses the registers with a 0x02 offset like 0x02, 0x06, 0x0A, 0x0E, ... For example, if the device is connected to SysData(31:16), then little endian software accesses the registers with a 0x02 offset with addresses like 0x02, 0x06, 0x0A, 0x0E, ... and big endian software accesses the registers with addresses like 0x00, 0x04, 0x08, 0x0C, ...

In the second option, the 16-bit device is halfword-aligned (using SysAddr bits n:1). With 16-bit systems it is imperative that the correct data line connections are made. Big Endian systems must attach SysData(31:16) and Little Endian systems must attach SysData(15:0). For example with half-word aligned connections, the software accesses the registers with addresses like 0x00, 0x02, 0x04, 0x06.

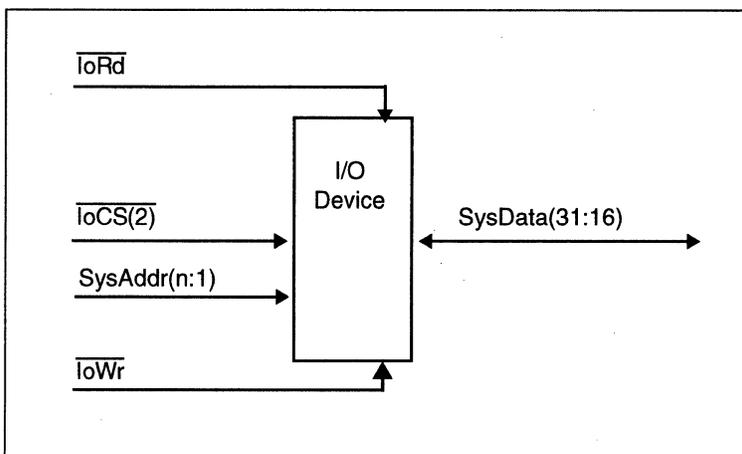


Figure 9.19 16-bit I/O System with Big Endian Connection

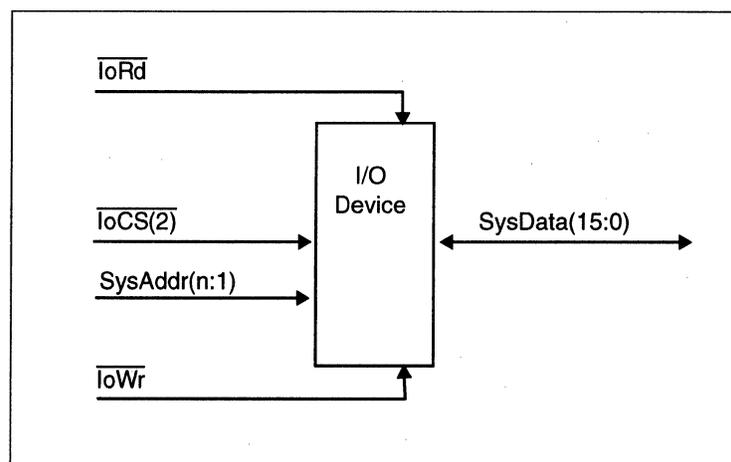


Figure 9.20 16-bit I/O System with Little Endian Connection

8-bit I/O Devices

Figure 9.21 and Figure 9.22 show two typical 8-bit I/O device systems. There are two choices for hooking up 8-bit I/O devices.

In choice 1, refer back to the 32-bit case in Figure 9.15 on page 23 and Figure 9.16 on page 23. The 8-bit device is word-aligned (using SysAddr bits $n:2$) even though it is a 8-bit device. This is the traditional MIPS connection and allows the device to be accessed from either endianness. For example, if the device is connected to SysData(7:0), then little endian software accesses the registers like 0x00, 0x04, 0x08, 0x0C, ... and big endian software accesses the registers with a 0x03 offset like 0x03, 0x07, 0x0B, 0x0F, ... For example, if the device is connected to SysData(31:24), then little endian software accesses the registers with a 0x03 offset with addresses like 0x03, 0x07, 0x0B, 0x0F, ... and big endian software accesses the registers with addresses like 0x00, 0x04, 0x08, 0x0C, ...

In choice 2, the 8-bit device is byte-aligned (using SysAddr bits $n:0$). With 8-bit systems it is imperative that the correct data line connections are made. Big Endian systems must attach SysData(31:24) and Little Endian systems must attach SysData(7:0). For example with half-word aligned connections, the software accesses the registers with addresses like 0x00, 0x01, 0x02, 0x03.

Since an 8-bit system is probably a smaller system, the SRAM/ROM can be attached directly to the SysAddr and SysData buses. In larger systems, FCT245 transceivers can be added between the memory bank and the SysData bus. Also in large systems, the SysAddr bus can also be buffered using FCT244 buffers.

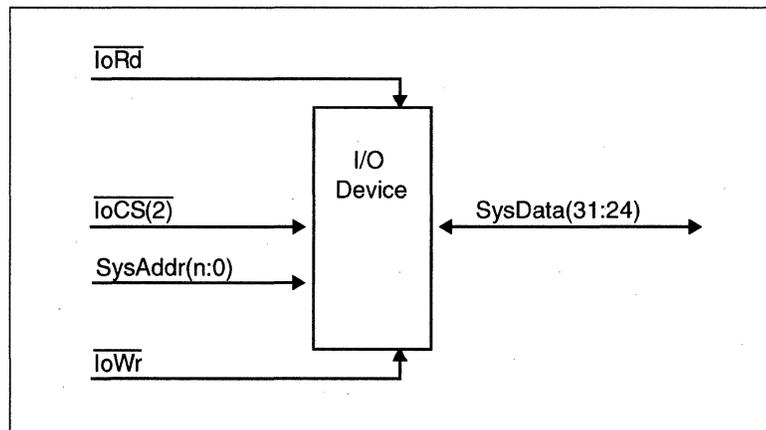


Figure 9.21 8-bit I/O System with Little Endian Connection

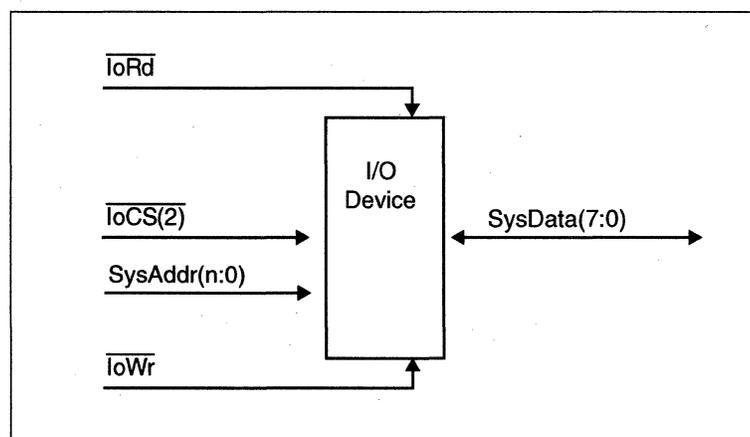
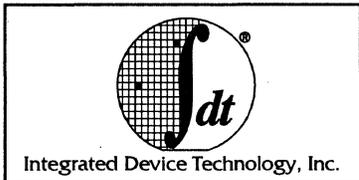


Figure 9.22 8-bit I/O System with Little Endian Connection

**Introduction**

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. One of the four on-chip bus controllers in the R36100 is the *DRAM Controller*.

This chapter provides an overview of the DRAM Controller interface, a complete description of the signal pins and their timing, and how the interface relates to typical external hardware DRAM systems.

Features

- Controls up to 4 banks of Page Mode DRAMs
- Each bank pair programmable to Interleaved or non-Interleaved mode
- Each bank programmable to use 1M, 4M, or 16M DRAM chips
- Each bank programmable to 32-bit or 16-bit Mode
- Provides jumper-less 16-bit to 32-bit or Interleaved upgrade
- Built-in $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ Refresh Timer
- Video DRAM Serial Transfer Protocol support
- Wait State Generator features:
 - Programmable time from start to end of each data access for each area
 - Programmable time options for Reads and Writes
 - Programmable time options for Single and Burst Accesses
 - Internally generates the RdCEN and AckN timing for all CPU accesses
 - A programmed value may be overridden by the $\overline{\text{SysWait}}$ input signal
- Direct Control of Data Path Transceivers include:
 - Direct Bus Connection
 - FCT260 Bidirectional Bus Exchanger Multiplexer
 - FCT245 Bidirectional Transceiver
 - FCT543 Bidirectional Registered Transceiver

Block Diagram

The functional block diagram of the DRAM Controller is shown in Figure 10.1. Located at the bottom of Figure 10.1, the DRAM Control Signal State Machine is responsible for generating the basic control signals used to connect to external DRAM chips and their transceivers. These signals include row and column address strobes, read enables, and write enables. The DRAM Controller as a whole works in cooperation with the Bus Interface Unit described in Chapter 7. Thus the Control Signal State Machine sends and receives information from the BIU Controller for assistance with controlling the port width and controlling partial word reads and writes.

The Control Signal State Machine also uses information stored in the software programmable DRAM Controller Register Bank for example, to control FCT260-Type versus FCT245-Type (transceiver interface) accesses. In addition to the Control Signal State Machine, there is a Refresh Timer and State Machine. The refresh circuitry implements $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh timing as required by conventional page mode DRAMs.

The DRAM Controller Wait-State Generator is shown in the center of Figure 10.1. The Wait-State Generator takes care of sending and receiving information from the BIU Controller in order to control the sequencing and timing of reading and writing each individual datum. The number of wait-states is derived from the settings programmed into the Register Bank. Once the correct number of wait-states has been counted out, the Wait-State Generator sets the appropriate internal BIU Acknowledge signals. The programmable Wait-State Generator eliminates the need for external state machines that are traditionally used for this function.

The DRAM Controller Decoder is shown at the top of Figure 10.1. The decoder constantly monitors the Bus Interface Unit's address and data bus to see if (1) the access is to the DRAM Controller's Register Bank, or, (2) the access is in one of the DRAM Controller's Chip Select areas that are responsible for controlling the bus transaction.

The DRAM Address Multiplexer is also shown at the top of Figure 10.1. The DRAM multiplexer switches the address lines between the MSB row address and LSB column address as required by conventional page mode DRAM chips. The multiplexer also includes address options to allow seamless upgrades from 16-bit to 32-bit or to interleaved 32-bit systems. The row address is also stored and compared using the Page Comparator circuitry. The Page Comparator allows the page mode DRAMs to enter into their faster page access mode whenever consecutive locations are accessed in the same block of memory.

The DRAM Controller Register Bank is shown at the left in Figure 10.1. The Register Bank allows the software programmer access to the many different options of the DRAM Controller. The chip select address ranges, the number of wait-states, the port-width of the chip select, and other similar options are programmed into the Register Bank as part of the software initialization sequence of the boot operating system.

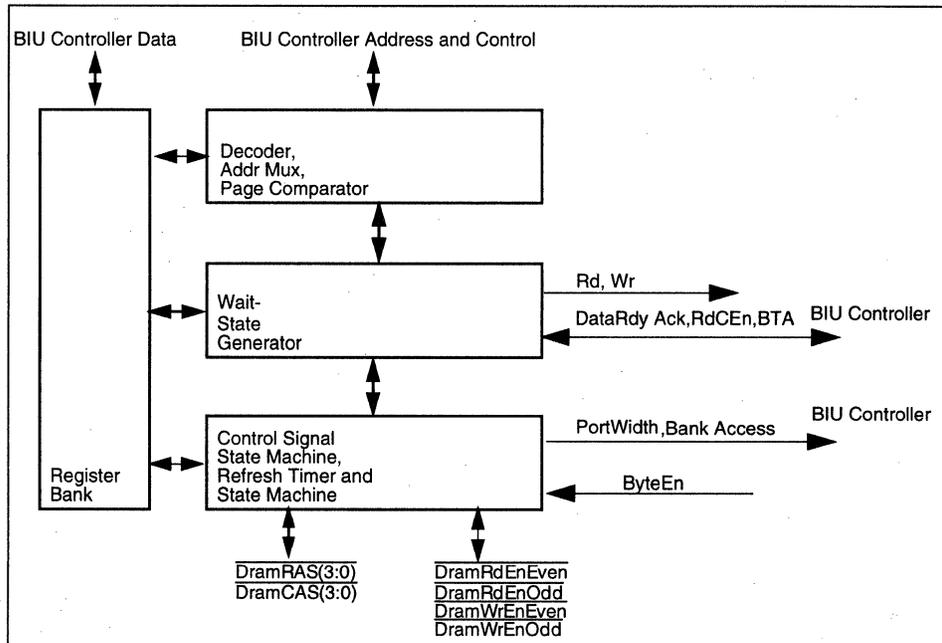


Figure 10.1 R36100 DRAM Bus Controller Block Diagram

DRAM BUS CONTROLLER INTERFACE SIGNALS

DRAM Interface Signals

The following external pins are typically attached directly from the R36100 RISController to external DRAM devices and their transceivers:

SysAddr(13:2) Output

The System Address provides the byte multiplexed address for DRAMs. This allows maximum of 16M words of unique locations to be accessed, thus providing a maximum of 64MBytes of memory for each bank. These signals share 12 of the lower 26 system address pins, SysAddr(13:2). Whenever a bus transaction is decoded to be a DRAM access, the behavior of these pins change and they act as DRAM-style multiplexed Row- and Column- address lines. Address assignments within the address multiplexer are such that a 16-bit system can be upgraded to a 32-bit system without external jumpers. In addition address assignments within the address multiplexer are such that single bank non-interleaved systems can be upgraded to pair-wise interleaved systems without external jumpers.

DramRAS(3:0) Output

DRAM Row Address Strobes are active low outputs used to strobe the row address into the DRAM. Each $\overline{\text{DramRAS}}(j)$ signal drives one bank of DRAM. They also provide balanced series resistor high drive for large memory systems (up to 8-10 chips each).

DramCAS(3:0) Output

DRAM Column Address Strobes are active low outputs used to strobe the column address into the DRAM. If the system uses a 16-bit wide bus instead of a 32-bit wide bus, then $\overline{\text{DramCAS}}(3:2)$ are used for a big endian system, while $\overline{\text{DramCAS}}(1:0)$ are used for a little endian system.

DramWrEnEven Output

DRAM Write Enable for Even Bank is an active low output signal used to write the selected DRAM bank 0 or 2. "Early write" cycles are used, so the byte selection is done by activating the leading edge of appropriate $\overline{\text{DramCAS}}(j)$ signals. It also provides balanced series resistor high drive for large memory systems (every other chip, up to 8-10 chips), although we recommend external FCT244/344 buffering if more than 8-10 connections are needed. Note that the DRAM specific write enables must be used instead of $\overline{\text{SysWr}}$ or $\overline{\text{MemWrEn}}(3:0)$ because refreshes may occur simultaneously with Memory Controller writes which could potentially cause 4M-16Mbit DRAMs to enter a test mode.

DramWrEnOdd Output

DRAM Write Enable for Odd Bank is an active low output signal used to write the selected DRAM bank 1 or 3. "Early write" cycles are used, so the byte selection is done by activating the leading edge of appropriate $\overline{\text{DramCAS}}(j)$ signals. It also provides balanced series resistor high drive for large memory systems (every other chip, up to 16 chips), although we recommend external FCT244/344 buffering if more than 8-10 connections are needed.

Note that the DRAM specific write enables must be used instead of $\overline{\text{SysWr}}$ or $\overline{\text{MemWrEn}}(3:0)$ because refreshes may occur simultaneously with Memory Controller writes, which could cause 4M-16Mbit DRAMs to enter a test mode.

DramRdEnEven Output

DRAM Read Enable for Even Bank is an active low output signal that is used to control the enabling of DRAM bank 0 or 2. Typically, DramRdEnEven is attached to the DRAM bank data transceiver output enable of banks 0 and 2, while DramCAS() controls the output enabling between the DRAM chips on the corresponding byte lane of bank 0 and 2.

In FCT260 type systems, DramRdEnEven is used as the overall DramRdEn path enable.

In FCT245 type systems, DramRdEnEven asserts on both reads and writes as a DramEnEven even bank transceiver enable.

DramRdEnOdd Output

DramTrEn

DRAM Read Enable for Odd Bank is an active low output signal used to control the enabling of DRAM bank 1 or 3. Typically DramRdEnOdd is attached to the DRAM bank data transceiver output enable of banks 1 and 3, while DramCAS() controls the output enabling between the two banks on each of the corresponding byte lanes. In FCT260-type systems, DramRdEnOdd is used as the overall DramRdPathSel path select. In FCT245-type systems, DramRdEnOdd asserts on both reads and writes as a DramEnOdd odd bank transceiver enable.

BIU Controller Signals

The BIU Controller Signals are used to complete the DRAM interface. These signals are also listed here for reference. Information specific to the DRAM Controller is given here and general information about the signal is given in Chapter 7, "System BIU Controller."

SysDataOutput/ Input

System Data Bus: A 32-bit peripheral connects the entire 32-bit SysData bus to its data pins or to its transceivers. 16-bit systems use the halfword offset address, A(x:1). (Note that the corresponding SysAddr() line is not SysAddr(x:1), since the DRAM address mux starts with SysAddr(2)). 8-bit systems can use the byte address, A(x:0). *In this case, 16-bit DRAMs connect to particular data pins depending on whether the Endianness of the system is Big Endian or Little Endian.* Thus 16-bit DRAMs use SysData(31:16) if they are Big Endian and SysData(15:0) if they are Little Endian. The User Mode Reverse Endianness Bit in the CP0 Status Register has no effect on the connections to SysData, however, it strongly recommended that the Reverse Endianness Bit not be used to "correct" an endianness connection as it does not function in the kernel mode address space.

SysWait Input

System Wait: The SysWait signal is ignored during DRAM Controller accesses.

Overview of the DRAM Controller

The R36100 RISController's DRAM Bus Controller supports up to four individual banks of standard RAS/CAS controlled page-mode DRAM chips. Each bank can have a minimum of 2x256KBytes (16-bit) and a maximum of 4x16MBytes (32-bit) of memory. Each bank can individually be programmed for 32-bits or 16-bits. Each pair of banks can be programmed to be non-interleaved or pair-wise interleaved. Thus the system as a whole can support up to four banks of DRAM and anywhere from 512KBytes up to 256MBytes.

The R36100 DRAM Controller, in addition to the $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ control lines, provides transceiver enable pins, and an address multiplexer (mux). The DRAM Controller also provides software configured options for wait-states as well as for CAS-before-RAS refresh timing.

Address mapping

All four banks must be contained within a single 256MByte address space. Then by programming the Base and Page Mask Registers of each bank, they can be individually mapped anywhere within the selected DRAM-designated 256MByte address space. This allows systems with mixed sized banks to have a contiguous memory space. Although mixed sized banks can be contained within a pair of DRAM banks, banks with different port widths (16-bit, 32-bit, or 64-bit) must be contained in different pairs of banks. Cacheability or non-cacheability of the references to these memory banks depends as to which virtual segment they are mapped to as per the R36100's memory map and CP0 Cache Control Register. This choice designates which references will be serviced as burst or non-burst references.

32-bit and 16-bit mode support

Each bank of DRAM memory can be programmed individually as either an interleaved, 32-bit, or a 16-bit bank. A two bit 'PortWidth' field is provided in the DRAM configuration register, one for each bank, which will select between the port width configuration options.

In 32-bit mode, the DRAM interface behaves similar to the 32-bit interface of the R3051 family in that address bits (1:0) are ignored as per word-aligned addressing. Thus address lines A(1) & A(0) are mapped out of the address multiplexing generation. Single word reads/writes are treated as one 32-bit datum. Partial word reads or writes are treated as single partial word read or write cycles by activating appropriate $\overline{\text{CAS}}$ signals, depending on the endianness of the processor. For block reads, the controller does four DRAM reads, back to back, to bring four words into the processor, using the page mode feature of DRAM.

The 16-bit mode is treated slightly different from the 32-bit mode. Most importantly, address line A(1) is included in the address multiplexer in order to support halfword-alignment. In case of 16-bit read/write instructions and data, the controller treats them as halfword mini-burst/burst read or write cycles, activating appropriate $\overline{\text{CAS}}$ signals based on the endianness of the device, $\overline{\text{DramCAS}}(3:2)$ for big endian and $\overline{\text{DramCAS}}(1:0)$ for little endian.

The mini-burst is continued until all halfword datum are read or written. Data will be driven on SysData(31:16) lines if it is a big-endian system or on SysData(15:0) if it is little-endian system. In cases of byte read or write instructions, the controller will activate appropriate $\overline{\text{CAS}}$ signals depending on the endianness of the processor.

With 16-bit mode block read accesses, the controller will perform 8 back to back reads in burst mode. This will be treated as a single burst transaction, bringing 16 bits of data every datum. The number of wait-states added between each datum can be programmed to differ from the number of wait-states added to the first datum, such that page mode DRAMs have optimal timing.

When the processor wants to read or write 32-bit data from a 16-bit memory, two 16-bit datum transfers occur back to back within the same address transaction. This is called a mini-burst write. The number of wait-states added between burst datum can be programmed to differ from the number of wait-states added to the first datum, such that page-mode DRAMs have optimal timing.

The interleaved mode is treated slightly different from the 32-bit mode. Most importantly, address line A(2) is ignored, along with A(1:0), per doubleword-aligned addressing. In essence, two words are simultaneously read from two separate banks; however, the second word must wait 1 extra clock before it can be latched into the CPU. Depending on the access time of the DRAM, this saves 1-2 clocks per every pair of words read. In addition, if every other word is latched, the CPU can pipeline the address for the second pair of words, a clock early, for additional savings.

Types of memory supported

The following three types of memory cycles are supported by the controller:

- **Page Mode support.** The page mode feature of the DRAM Controller is always enabled. In the case of a mini-burst or burst refill, the page mode is used to obtain data by use of an octi (16-bit) or quad (32-bit) word read. In the cases specified in a later section, "Page Comparator Algorithm", \overline{RAS} will be left active expecting a subsequent page mode access to the same block of memory. The controller has an on-chip page register and comparator which uses the programmed DRAM density to determine whether or not a given access can take advantage of page mode; as well as whether or not to leave \overline{RAS} asserted at the end of the transaction.
- **Non-Interleaved support.** At any given time in the non-interleaved mode, only one bank will have an active \overline{RAS} . In the case of an access to a different bank, first the \overline{RAS} of the active bank will be deactivated and then the \overline{RAS} of the accessed bank will be activated. In case of a page miss within the same bank, \overline{RAS} will deactivate for a precharge period, a new page address will be strobed in by driving the new row address, and then re-activating \overline{RAS} .
- **Interleaved support.** A programmable option field is provided in the configuration register which will enable or disable two-way interleaving. Various interleaving sub-options allow various types of transceivers to be used by changing the functionality of the controller's output enables.

Programmable wait state generation

A programmable wait state feature is supported by:

- \overline{RAS} Precharge, Row Addr Setup, and Row Addr Hold settings
- \overline{CAS} Addr Setup/Precharge setting
- \overline{CAS} Pulse Width and \overline{CAS} Addr Hold settings

Page Comparator Algorithm

An internal page comparator compares the page address of consecutive DRAM bus cycles. After completing the current DRAM bus access, \overline{RAS} can be held asserted after:

- Writes
- Single Word Reads
- Burst Cache Reads

When asserted, \overline{RAS} is de-asserted if:

- refresh occurs
- non-page write occurs
- non-page read occurs

The page comparator is not affected by non-DRAM accesses. It is assumed that uncached reads and writes are unlikely to be done to DRAM, thus the distinction between instruction and data is not statistically important to throughput. Also, the maximum assertion time for \overline{RAS} is assumed to be covered by the occurrence of refreshes.

Unaligned page accesses

Since long bursts are always aligned to the burst length, bursts across non-page boundaries are not possible from the CPU or internal DMA channels. Therefore, unaligned page accesses will not occur. There is one possible exception: if external DMA with long burst access does a non-aligned burst, then a page boundary crossing is possible. Most DMA agents capable of long bursts (for example, 16 words) also pre-align or are capable of pre-aligning the burst on a boundary (for example, align to a 16 word boundary).

Refresh Timing

The $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh mode is supported. The refresh rate is programmable in order to take into account the speed of the processor.

Initialization

The system boot software is responsible for initializing the DRAMs after reset. The DRAM Controller is guaranteed to hold all DRAM control signals de-asserted until a proper DRAM cycle is initiated by the user. Usually initialization involves the software OS to program a wait of 200us after power up (reset), initializing all of the DRAM control registers, and then doing 8 $\overline{\text{RAS}}$ cycles. Alternatively, the boot software can wait 200us and then wait until 8 refresh cycles occur.

Programmable features

The DRAM Controller has the following programmable features:

- Page Size
- $\overline{\text{RAS}}$ assertion selection
- $\overline{\text{RAS}}$ precharge time
- $\overline{\text{RAS}}$ Addr Setup and Addr Hold time
- $\overline{\text{CAS}}$ precharge/Addr Setup time
- $\overline{\text{CAS}}$ Addr Hold Time on Writes (WrBTA)
- $\overline{\text{CAS}}$ Pulse Width
- Internal Burst Ack generation

Signal Control Interface

The DRAM Controller provides read enables and write enables that are suitable for direct chip connection. The read and write enables can in general also be attached to FCT260, FCT245, and FCT543 type transceivers.

Wait State Generator

The Wait-State Generator controls the speed of the DRAM accesses to and from the Bus Interface Unit Controller. This includes the time from the start of a DRAM transaction until the first datum is sent or received. The Wait-State Generator also is programmed to generate the internal RdCEnN and AckN signals for CPU read and write requests.

The internal Acknowledge signal, "AckN" (as described in Chapter 7), is the same as the external signal pin that the R3051 RISController family uses. On single word reads and on both single word and burst writes, AckN is automatically placed at the end of the transaction by the Wait-State Generator. Burst DRAM read operations return AckN earlier than the end of the transaction (because of the read buffer); thus, a Control Register'Burst Ack' field is provided.

Register Option Field Programmability

The DRAM Controller contains 4 sets of registers, one set for each chip select, $\overline{\text{DramRAS}}(3:0)$. There is also a global set of registers for the address multiplexer options and refresh timing options. These registers allow the DRAM Controller to be configured for different speeds and types of DRAM chips; therefore, almost any system speed/cost/manufacturing trade-off can be accommodated.

Register Descriptions

Table 10.1 provides the address map for the DRAM Controller registers. the Note that Big Endian software must offset these addresses by b'10 (0x2).

Phys. Address	Register
0xFFFF_E100 0xFFFF_E104	DRAM Refresh Count Register DRAM Refresh Compare Register
0xFFFF_E110 0xFFFF_E114	DRAM <u>RAS</u> Multiplexer Select Register for Pair 1:0 DRAM <u>RAS</u> Multiplexer Select Register for Pair 3:2
0xFFFF_E120 0xFFFF_E124	DRAM <u>CAS</u> Multiplexer Select Register Pair 1:0 DRAM <u>CAS</u> Multiplexer Select Register Pair 3:2
0xFFFF_E180 0xFFFF_E184 0xFFFF_E188 0xFFFF_E18C	DRAM MSB Base Address Register for Bank 0 DRAM MSB Bank Mask Register for Bank 0 DRAM LSB Control Register for Bank 0 DRAM MSB Control Register for Bank 0
0xFFFF_E190 0xFFFF_E194 0xFFFF_E198 0xFFFF_E19C	DRAM MSB Base Address Register for Bank 1 DRAM MSB Bank Mask Register for Bank 1 DRAM LSB Control Register for Bank 1 DRAM MSB Control Register for Bank 1
0xFFFF_E1A0 0xFFFF_E1A4 0xFFFF_E1A8 0xFFFF_E1AC	DRAM MSB Base Address Register for Bank 2 DRAM MSB Bank Mask Register for Bank 2 DRAM LSB Control Register for Bank 2 DRAM LSB Control Register for Bank 2
0xFFFF_E1B0 0xFFFF_E1B4 0xFFFF_E1B8 0xFFFF_E1BC	DRAM MSB Base Address Register for Bank 3 DRAM MSB Bank Mask Register for Bank 3 DRAM LSB Control Register for Bank 3 DRAM MSB Control Register for Bank 3

Table 10.1 DRAM Controller Registers.

DRAM Refresh Count Register ('DramRefreshCountReg')

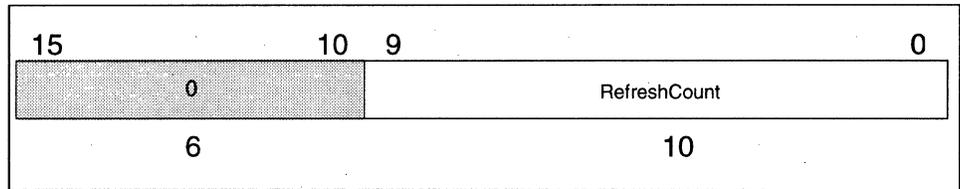


Figure 10.2 DRAM Refresh Count Register ('DramRefreshCountReg')

BIT	Function
9:0	RefreshCount

Table 10.2 DRAM Refresh Count Register (DramRefreshCountReg) Bit Assignments.

The lower 10 bits form a 10-bit binary up-counter. The Count register, shown in Figure 10.2, ticks upward on each system clock. When Count equals Compare, the DRAM Controller will initiate a refresh sequence and the Count register will be reset back to 0. The upper 6 bits are reserved to be "0". The default value of the DRAM Refresh Count Register, shown in Table 10.2, is 0x0000 at reset. The register is both readable and writable.

Staggered Refresh

In order to reduce the amount of peak instantaneous current and intra-bus transaction average current used by refreshing DRAMs, refresh is done by refreshing (RAS'ing) Banks 0 & 2 together, then afterwards refreshing (RAS'ing) Banks 1 & 3 together.

Refresh Arbitration

Refreshes on the R36100 must obtain the DRAM Controller before doing a refresh. DRAM systems in general must not use $\overline{\text{MemWrEn}}$ or $\overline{\text{SysWr}}$, because some other peripheral driving them low during a refresh would accidentally put some types of DRAM chips in their internal test mode. If the CPU or DMA channel tries to access DRAM at the same time as a Refresh, they will wait for the Refresh to finish.

Panic Mode Refresh Application

Ordinarily it is only possible to have a 4-word burst DMA after which the refresh controller can regain the bus and do a pending refresh. However, external DMA can burst up to a system defined length. In such a situation, multiple (depending on the application) refresh ticks may be missed. If this is of concern to the system designer, they can either (1) divide the external DMA burst into smaller units, or, (2) initiate N refreshes before and N refreshes after the burst, where N is the potential number of refreshes missed.

Reduced Frequency Mode Application

To use DRAMs in the reduced frequency mode, it is assumed the ultimate objective is to save power. Suggestion 1 is to use self-refreshing DRAMs. Suggestion 2 is to use low power extended refresh DRAMs. Since the indirect objective is to minimize the $\overline{\text{RAS}}$ low time, a $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ Refresh should complete as soon as possible. Therefore, one way of accomplishing this is to reprogram the Refresh count/compare registers to suitable values such that once every refresh period (64ms) the CPU is internally interrupted. The interrupt will exit the halt and RF modes. A short interrupt handler loop can strobe through all 512 row addresses and then return the CPU into halt and RF modes.

DRAM Refresh Compare Register ('DramRefreshCompReg')

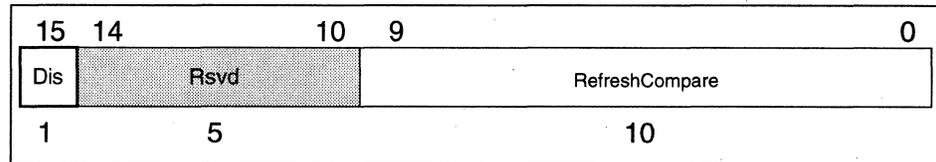


Figure 10.3 DRAM Refresh Compare Register.

This register forms a 10-bit Compare Register, shown in Figure 10.3. Bit 15 is a Disable Field, and bits 14:10 are reserved and should be written to with the same value as that of bit 15. When Compare equals Count, the DRAM Controller will reset the count back to 0. If the Refresh Disable Field is set to 'enable', then a refresh sequence is initiated. The default value of the DRAM Refresh Count Register is 0x0000 at reset. As an example: for 25 MHz CPU with 8ms/512 refresh period, Compare should be programmed to $\text{Floor}(8\text{m}/(512+1) / (1/25\text{M})) - 1 = 0\text{x}0185$. The register is both readable and writable. Table 10.3 lists the bits assignments for this register.

Note: Technical worst case accounts for maximum burst length, where it is sufficient to add 1 to the DRAM page size.

The Refresh Compare Register is provided in a binary count manner (as opposed to a frequency select manner) to allow easier accessibility for diagnostic and test purposes. The default value at reset is 0xFFFF (refer to Table 10.4). Common refresh settings are given in Table 10.5.

Bit	Function
15	Refresh Disable
14:10	Reserved (write the same value as bit 15)
9:0	Refresh Count

Table 10.3 DRAM Refresh Compare Register ('DramRefreshCompareReg') Bit Assignments.

Value	Action
'1'	Disable Refresh Counter (0xFFFF default)
'0'	Enable Refresh Counter

Table 10.4 Refresh Disable ('RefreshDis') Field Encodings.

Refresh Compare	CPU Frequency
0x00F8	16
0x0136	20
0x0184	25
0x0201	33
0xFFFF	default: disabled

Table 10.5 Common Refresh Settings for 8ms/512 or 16ms/1024 DRAMs.

DRAM $\overline{\text{RAS}}$ Multiplexer Select Register for Pair(1:0, 3:2)
 ('DramRasMuxSelReg'1_0, 3_2)

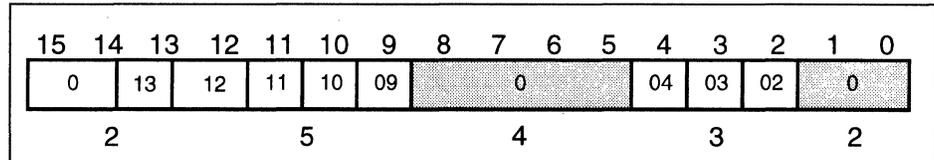


Figure 10.4 DRAM $\overline{\text{RAS}}$ Mux Select Register ('DramRasMuxSelReg').

The DRAM $\overline{\text{RAS}}$ Address Multiplexer Select Register, shown in Figure 10.4, programs which address bits go out to a DRAM Pair system during the row address period. The different selections allow software to upgrade the size of the DRAM chips and the memory port width without the use of external hardware jumpers. The register is both readable and writable with a default of 0x0000 at reset. This register should be programmed before the DRAM Controller is first used.

Bit	Function	High	Low
13	SysAddr RAS_13	A26	A24
12	SysAddr RAS 12	A23	A10
11	SysAddr RAS 11	A22	A20
10	SysAddr RAS 10	A19	A10
09	SysAddr RAS 09	A18	A9
04	SysAddr RAS 04	A25	A13
03	SysAddr RAS 03	A22	A12
02	SysAddr RAS 02	A21	A11

Table 10.6 DRAM $\overline{\text{RAS}}$ Mux Select Register Bit Assignments.

Note: Bits 15, 14, 8:5, 1, and 0 are reserved for future use.

DRAM $\overline{\text{CAS}}$ Multiplexer Select Register for Pair (1:0, 3:2)
 ('DramCasMuxSelReg'1_0, 3_2)

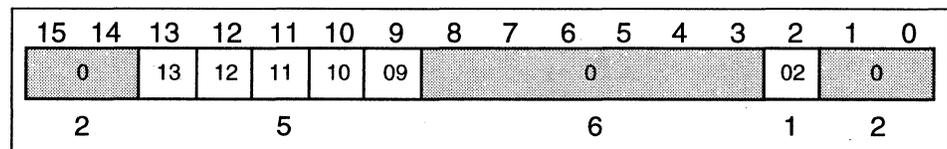


Figure 10.5 DRAM CAS Mux Select Register

The DRAM Pair $\overline{\text{CAS}}$ Multiplexer Select register, shown in Figure 10.5, programs which address bits go out to the DRAM system during the column address period. The different selections allow software to upgrade the size of the DRAM chips and the memory port width without the use of external hardware jumpers. The register is both readable and writable with a default of 0x0000 at reset. This register should be programmed before the DRAM Controller is first used. The DRAM $\overline{\text{CAS}}$ Mux Select Register Bit Assignments are listed in Table 10.7. Refer to Table 10.8 for an example of DRAM $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ Mux select register settings.

Bit	Function	High	Low
13	SysAddr CAS 13	A13	A1
12	SysAddr CAS 12	A24	A12
11	SysAddr CAS 11	A11	A1
10	SysAddr CAS 10	A10	A1
09	SysAddr CAS 09	A9	A0
02	SysAddr CAS 02	A20	A2



Table 10.7 DRAM $\overline{\text{CAS}}$ Mux Select Register ('DramCasMuxSelReg') Bit Assignments.

Note: Bits 15, 14, 8:3, 1, and 0 are reserved for future use.

SysAddr	32-bit 256K		32-bit 1M		32-bit 4M		Row 16-bit	Col 256K	Row 16-bit	Col 1M	Row 16-bit	Col 4M
	Row 32-bit	Col 256K	Row 32-bit	Col 1M	Row 32-bit	Col 4M						
SysAddr(13)												
SysAddr(12)					A23	A12					A10	A12
SysAddr(11)			A20	A11	A20	A11			A20	A1	A20	A11
SysAddr(10)	A19	A10	A19	A10	A19	A10	A10	A1	A19	A10	A19	A1
SysAddr(9)	A18	A9	A18	A9	A18	A9	A18	A9	A18	A9	A18	A9
SysAddr(8)	A17	A8	A17	A8	A17	A8	A17	A8	A17	A8	A17	A8
SysAddr(7)	A16	A7	A16	A7	A16	A7	A16	A7	A16	A7	A16	A7
SysAddr(6)	A15	A6	A15	A6	A15	A6	A15	A6	A15	A6	A15	A6
SysAddr(5)	A14	A5	A14	A5	A14	A5	A14	A5	A14	A5	A14	A5
SysAddr(4)	A13	A4	A13	A4	A13	A4	A13	A4	A13	A4	A13	A4
SysAddr(3)	A12	A3	A12	A3	A22	A3	A12	A3	A12	A3	A22	A3
SysAddr(2)	A11	A2	A21	A2	A21	A2	A11	A2	A11	A2	A21	A2
SysAddr(1)												
SysAddr(0)												

SysAddr	Row 64-bit	Col 256K	Row 64-bit	Col 1M	Row 64-bit	Col 4M
SysAddr(13)						
SysAddr(12)					A23	A24
SysAddr(11)				A22	A11	A22
SysAddr(10)	A19	A10	A19	A10	A19	A10
SysAddr(9)	A18	A9	A18	A9	A18	A9
SysAddr(8)	A17	A8	A17	A8	A17	A8
SysAddr(7)	A16	A7	A16	A7	A16	A7
SysAddr(6)	A15	A6	A15	A6	A15	A6
SysAddr(5)	A14	A5	A14	A5	A14	A5
SysAddr(4)	A13	A4	A13	A4	A13	A4
SysAddr(3)	A12	A3	A12	A3	A12	A3
SysAddr(2)	A11	A20	A21	A20	A21	A20
SysAddr(1)						
SysAddr(0)						

SysAddr	Row 16-bit	Col 16M	Row 2-bit	Col 16M	Row 64-bit	Col 16M
SysAddr(13)	A24	A1	A24	A13	A26	A13
SysAddr(12)	A23	A12	A23	A12	A23	A24
SysAddr(11)	A20	A11	A20	A11	A22	A11
SysAddr(10)	A19	A10	A19	A10	A19	A10
SysAddr(9)	A18	A9	A18	A9	A18	A9
SysAddr(8)	A17	A8	A17	A8	A17	A8
SysAddr(7)	A16	A7	A16	A7	A16	A7
SysAddr(6)	A15	A6	A15	A6	A15	A6
SysAddr(5)	A14	A5	A14	A5	A14	A5
SysAddr(4)	A13	A4	A25	A4	A25	A4
SysAddr(3)	A22	A3	A22	A3	A12	A3
SysAddr(2)	A21	A2	A21	A2	A21	A20
SysAddr(1)						
SysAddr(0)						

Table 10.8 Example 'DramRasMuxSelReg' and 'Dram-CasMuxSelReg' Settings.

DRAM MSB Base Address Register for Bank 0..3
 ('DramMSBBaseAddrReg(0..3)')

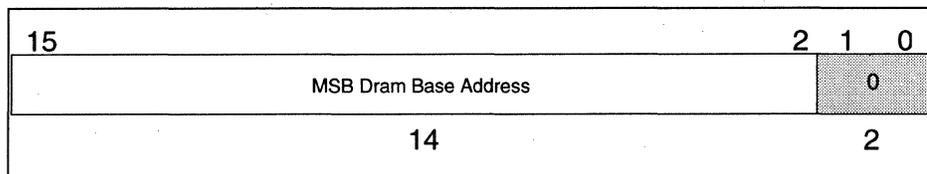


Figure 10.6 DRAM MSB Base Address Register ("DramMSBBaseAddrReg")

This field contains Bits 31-18 of the starting base physical address of the DRAM Bank. The programmer must write the same value for bits 31-28 to all DRAM MSB Base Address Registers. In addition, the programmer must write "0" for bits 17-16 of all DRAM MSB Base Address Registers, which restricts the smallest multiple DRAM bank size to 256K. The default value on reset is 0xEEEE. Thus the upper 3 DRAM Bank MSB Base Address registers must be programmed before any DRAM accesses can be initiated. The register is both readable and writable. Figure 10.6 illustrates the DRAM MSB Base Address Register.

Internally to the R36100, bits 31-28 of Bank 0 MSB Base Address are used for the starting address of all banks. Also internally, bits 17-16 are reserved and hardwired to 0.

An example for 4 banks (2 pairs) of 1MByte interleaved DRAM starting at physical address 0 is in Table 10.9.

Bank	Phys. Address
Bank0	0x0000_0000
Bank1	0x0000_0000
Bank2	0x0020_0000
Bank3	0x0020_0000

Table 10.9 Example Bank Base Address Register ('DramMSBBaseAddrReg') Assignment.

An example for 1M DRAM + 2 banks of 4MByte interleaved DRAM at physical address 0 is in Table 10.10. Note that Bank1 must be assigned to an unused memory space.

Bank	Phys. Address
Bank0	0x0000_0000
Bank1	0x0F00_0000
Bank2	0x0010_0000
Bank3	0x0010_0000

Table 10.10 Example Bank Base Address Register ('DramMSBBaseAddrReg') Assignment.

**DRAM MSB Bank Mask Register for Bank 0..3
(‘DramMSBBankMaskReg(0..3)’)**

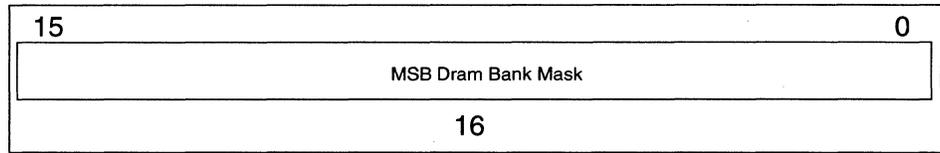


Figure 10.7 DRAM MSB Bank Mask (‘DramMSBBankMask(3:0)’) Registers.

There are 4 bank mask registers, one for each DRAM bank. The bank mask address register, shown in Figure 10.7, represents the most significant 16 address bits (bits 31:16). Bit settings for this register are listed in Table 10.11.

The bank mask registers are used to decide which address bits in the base address are to be used for comparing whether a DRAM bank select is to be activated. This DRAM Bank Mask is independent of the DRAM $\overline{\text{RAS}}$ Page Size Mask.

Internally, bits 31-16 must be programmed to the desired bank mask. This corresponds to separate address spaces for each chip select of 64K to 256M. Internally, bits 15:0 are ignored for bank mask comparisons.

To summarize, Bits 31:16 of each DRAM bank page mask are used to distinguish the size of each memory space. The format of the DramMSB-BankMask is displayed in the above figure. The register is both readable and writable and is set to 0xFFFF by default on reset. This register should be programmed before the DRAM Controller is first used.

Value	Action
'1'	Bit is used in comparison
'0'	Bit is masked out of comparison

Table 10.11 DRAM MSB Bank Mask Bit Settings.

**DRAM LSB Control Register for Bank 0..3
(‘DramLSBControlReg(0..3)’)**

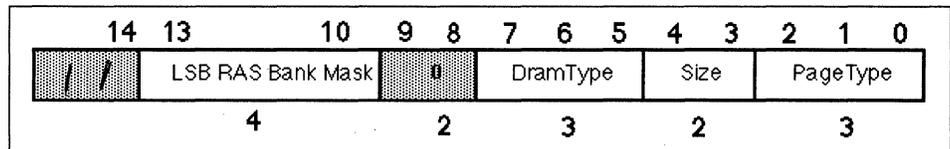


Figure 10.8 DRAM LSB Bank Control Register (‘DramLSBControlReg’).

The DRAM LSB Control Register, shown in Figure 10.8, is used to control various DRAM controller options. The register is both readable and writable. The default value at reset is 0xFC03. This register should be programmed before the DRAM Controller is first used. Bit assignments for this register are listed in Table 10.12.

Bit	Description
15:14	Reserved to 1
13:10	RASBankMask
9:8	Reserved to 0
7:5	DramType
4:3	Size
2:0	PageType

Table 10.12 DRAM LSB Control Register (‘DramLSBControlReg’) Bit Assignments.

RAS Bank Mask ('RASPageMask') Field (bits 15:8)

The RAS Bank Compare Mask is used to determine how many of the upper physical address bits will be compared to determine if subsequent DRAM accesses are on the same RAS page of memory and thus do not need to initiate a RAS precharge. Note that the RAS Bank Mask is independent and does not have the same function as the DRAM Bank Mask. To determine the bank mask, Physical Address Bits 31-16 are always compared, and Physical Address Bits 15-8 are compared if their corresponding RAS Bank Mask bit is clear. Page Mask bits are listed in Table 10.13 and DRAM LSB Page Mask bit settings are listed in Table 10.14.

DRAM-Type	Port-Size	Inter-leaved(1)	PageMask to LSB(2)
16MxN	64-bit	yes	11
16MxN	32-bit	non-	13(2)
16MxN	16-bit	non-	12
4MxN	64-bit	yes	11
4MxN	32-bit	non-	12
4MxN	16-bit	non-	9
1MxN	64-bit	yes	11
1MxN	32-bit	non-	11
1MxN	16-bit	non-	10
256KxN	64-bit	yes	10
256KxN	32-bit	non-	10
256KxN	16-bit	non-	9

NOTES:
 1. Interleaved systems compare 1 less bit than theoretically possible, due to Column Address selection limitations.
 2. Most configurations compare less bits than theoretically possible, as a trade-off for the jumper-less expansion.
 3. 16Mx32 systems compare less bits than theoretically possible, due to Column Address selection limitations.
 4. For very small memory systems, don't set any of the Page Type ('PType') control bits. In this case, the page comparator is ignored.

Table 10.13 Page Mask ('PMask') Bits.

Value	Action
'1'	Bit is used in comparison
'0'	Bit is masked out of comparison

Table 10.14 DRAM LSB Page Mask Bit Settings.

DRAM Type ('DramType') Field (bits 7:5)

The DRAM Type selections are used to chose between different types of DRAM configurations. DRAM Type settings are listed in Table 10.15.

Value	Action
'7:3'	Reserved
'2'	FCT260
'1'	FCT245
'0'	FCT543 (default)

Table 10.15 DRAM Type ('DramType') Settings.

FCT543-Type (Latched Non-Multiplexer Type)

FCT543 Latched Mode assumes that the transceiver hardware between DRAM and the R36100 consists of latched or registered transceivers such as the FCT543.

Setting this type with an interleaved port width causes the DRAM controller to use the interleave bus protocol. Interleaved systems use a different bus protocol which in essence accesses two banks at the same time, but only output enables 1 bank at a time, and thus can be burst read very quickly.

Non-interleaved systems use a bus protocol which in essence accesses the two banks in separate distinct address spaces. Latching words 1 and 3 on burst reads allows words 1 and 3 to be returned exactly 1 clock after words 0 and 2, respectively. This in turn allows the LS double word address to be bumped up (from 0x00 to 0x08) one clock earlier.

FCT245 Type (Non-latched Transceiver Type)

Non-Latched Mode assumes that the transceiver hardware between DRAM and the R36100 consists of non-latched or non-registered transceivers such as the FCT245. In addition, in order to support the direction select and the single output enables of FCT245s, the even read output enable, $\overline{\text{DramRdOEnEven}}$, logically OR's $\overline{\text{DramRdEnEven}}$ with $\overline{\text{DramWrEnEven}}$. Likewise, the odd read output enable, $\overline{\text{DramRdOEnOdd}}$, logically OR's $\overline{\text{DramRdEnOdd}}$ with $\overline{\text{DramWrEnOdd}}$. $\overline{\text{DramWrEnEven}}$ and $\overline{\text{DramWrEnOdd}}$ are unchanged as they are needed to indicate writes. If the FCT245 Type is used, all four banks must be of FCT245 Type.

FCT260-Type (Latched Multiplexer Type)

FCT260 Latched Multiplexer Type assumes that the transceiver hardware between DRAM and the R36100 consists of latched or registered multiplexers such as the FCT260. In order to support the path select and single output enables of multiplexers, the even read output enable, $\overline{\text{DramRdOEnEven}}$, logically OR's $\overline{\text{DramRdOEnEven}}$ with $\overline{\text{DramRdOEnOdd}}$ and stays low for most of the bus transaction instead of toggling just for even banks, so that it can be used for the FCT260 output enable. $\overline{\text{DramRdOEnOdd}}$ is unchanged, so that it can be used for a path select.

Setting this type with an interleaved port width causes the DRAM controller to use the interleave bus protocol. Interleaved systems use a different bus protocol which in essence accesses two banks at the same time, but only enables 1 bank at a time, and thus can be burst read very quickly. Non-interleaved systems use a bus protocol which in essence accesses the two banks in separate distinct address spaces. Latching words 1 and 3 on burst reads allows words 1 and 3 to be returned exactly 1 clock after words 0 and 2, respectively. This in turn allows the LS double word address to be bumped up (from 0x00 to 0x08) one clock earlier.

Note: At present, the FCT260 hardware approach represents one of the better price/performance ratios for interleaved systems, since only 3 chips instead of 4 chips are required.

Port Size ('Size') Field (bits 4:3)

The Port size of a bank determines its memory width. Table 10.16 lists the DRAM Port Width Encoding field.

Value	Port Size
'10'	16-bit
'00'	32-bit (default)
'11'	64-bit (2x32-bit interleaved)
'01'	Reserved

Table 10.16 DRAM Port Width ('Size') Encoding Field.

Page Type ('PageType') Field (bits 2:0)

The Page Type field determines after which type of bus transaction $\text{DramRAS}\bar{0}$ is held low. The values and descriptions for this field are indicated in Table 10.17. **Note:** Default value is '111'.

Value	Description
'1xx'	After Burst Reads.
'0xx'	Not After Burst Reads.
'x1x'	After Single Word Reads.
'x0x'	Not After Single Word Reads.
'xx1'	After Writes.
'xx0'	Not after Writes.

Table 10.17 Page Type ('PageType') (bits 2:0) Field.

DRAM MSB Control Register for Bank 0..3 ('DramMSBControlReg'0..3)

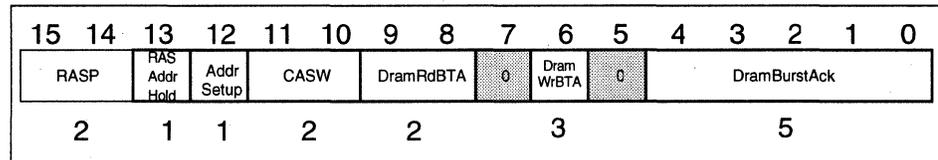


Figure 10.9 DRAM MSB Bank Control Register ('DramMSBControlReg').

The DRAM MSB Control Register, shown in Figure 10.9, is used to control various DRAM controller options. This register is both readable and writable. Bit assignments for this register are listed in Table 10.18.

Bit	Description
15:14	RASP
13	RASAddrHold
12	AddrSetup
11:10	CASW
9:8	DramRdBTA
7	'0'
6	DramWrBTA
5	'0'
4:0	DramBurstAck

Table 10.18 DRAM MSB Control Register Bit Assignments.

RAS Precharge Period ('RASP') Field (bit 15:14)

Before initiating a DRAM access to a new page, $\overline{\text{RAS}}$ must be held de-asserted in order to precharge the DRAM chip row decoders and sense amps. The RASP setting defines the length of this precharge period. The default value at reset is '0' which encodes to 1 clock. RAS Precharge Field Encodings are listed in Table 10.19.

Value	Action
'3'	reserved
'2'	3 clocks
'1'	2 clocks
'0'	1 clock (default)

Table 10.19 RAS Precharge ('RASP') Field Encodings.

RAS Address Hold Time ('RASAddrHold') Field (bit 13)

DRAM Address Hold Time is required in the following three places:

- after $\overline{\text{RAS}}$ asserts
- after $\overline{\text{CAS}}$ asserts
- after $\overline{\text{CAS}}$ re-asserts

Address Hold Time from $\overline{\text{RAS}}$ asserting is handled by this field, the RASAddrHold field. Address Hold Time from $\overline{\text{CAS}}$ asserting and re-asserting on reads is handled by the CASW field. Address Hold Time from $\overline{\text{CAS}}$ asserting and re-asserting (on writes) is handled by the DRAMWrBTA field.

RASAddrHold (see Table 10.20 for field encodings) defines the length of the DRAM row address hold time. Normally, 0.5 clocks is enough hold time since most DRAMs require that the row address be held for about 10ns after $\overline{\text{RAS}}$ asserts. However, in very fast systems where the clock period is short, or in very noisy, heavily delayed systems, additional address hold time may be needed. Thus RASAddrHold can be extended from 0.5 to 1.5 clocks if necessary.

Value	Action
'1'	1.5 clocks
'0'	0.5 clocks (default)

Table 10.20 RAS Address Hold Time ('RASAddrHold') Field Encoding.

Address Setup Time to $\overline{\text{RAS}}$ and to $\overline{\text{CAS}}$ ('AddrSetup') Field (bit 12)

DRAM Address Setup Time is required in the following three places:

- Row Address Setup Time to $\overline{\text{RAS}}$ asserting
- Column Address Setup Time to $\overline{\text{CAS}}$ asserting (also Early Write Signal Setup Time to $\overline{\text{CAS}}$ asserting)
- Column Address Setup Time to $\overline{\text{CAS}}$ re-asserting on mini-bursts or bursts (also Early Write Signal Setup Time to $\overline{\text{CAS}}$ re-asserting on mini-bursts or bursts)

AddrSetup (field encodings are listed in Table 10.21) defines the length of the DRAM address setup time. Normally, 0.5 clocks is enough setup time since most DRAMs require that the row or column address be setup 0ns before $\overline{\text{RAS}}$ or $\overline{\text{CAS}}$ asserts. However, in very fast systems where the clock period is short, or in very noisy, heavily delayed systems additional address (and also early write) setup time may be needed. Therefore, AddrSetup can be extended from 0.5 to 1.5 clocks if necessary. Thus for a new page DRAM access, AddrSetup may add 1 extra address setup clock cycle before $\overline{\text{RAS}}$ asserts, before $\overline{\text{CAS}}$ asserts, or before every $\overline{\text{CAS}}$ re-assertion.

Value	Action
'1'	1.5 clocks
'0'	0.5 clocks (default)

Table 10.21 Address Setup Time to $\overline{\text{RAS}}$ or to $\overline{\text{CAS}}$ ('AddrSetup') Field Encoding.

CAS Active Pulse Width ('CASW') Field (bit 11:10)

The 2 bit encoding lengthens the $\overline{\text{CAS}}$ active pulse width from a minimum of 0.5 clocks up to 2.5 clocks. The default value is '1' at reset which decodes to 1.5 clocks (see Table 10.22 for field encodings). This field can also be thought of as the CASAddrHold field on read accesses; however, on DRAM writes, DRAMWrBTA must be used to extend the address hold and early write signal hold time.

Value	Action
'3'	reserved
'2'	2.5 clocks
'1'	1.5 clocks (default)
'0'	0.5 clocks

Table 10.22 CAS Width ('CASW') Field Encoding.

DRAM Read Cycle Bus Turn-Around ('DramRdBTA') Field (bit 9:8)

The Bus Turn-Around field determines the minimum number of clocks between the end of a read and the beginning of the next non-DRAM bus transaction. Sometimes a slow interface is needed because of the amount of time it takes the DRAM or its transceivers to tri-state off of their respective busses. Thus, a slow bus turnaround option is incorporated into the DRAM Controller. A two-bit value stored in a control register will stall the bus interface unit between the end of a read cycle and from starting the subsequent transfer by up to two system clock cycles. The default at reset is the value '1', which decodes to 1 clock of BTA. Field encodings are listed in Table 10.23.

Value	Action
'3'	3 clocks
'2'	2 clocks
'1'	1 clock (default)
'0'	0 clocks

Table 10.23 DRAM Read Cycle Bus Turn-Around ('DramRdBTA') Field Encoding.

DRAM Write Cycle Bus Turn Around ('DramWrBTA') Field (bit 6)

The DRAM Write Cycle Bus Turn Around Field determines whether other transactions can begin either one clock cycle before the final $\overline{\text{CAS}}$ de-assertion or after the final $\overline{\text{CAS}}$ de-assertion.

DramWrBTA defines the minimum number of clocks between the end of a write and the beginning of the next non-DRAM bus transaction. The DRAM Controller uses the early write protocol and thus can typically give up the data and address buses one clock before the end of the write (1 clock before $\overline{\text{CAS}}$ de-asserts for the last time).

However, this leaves the column address and the early write signal 0.5 clocks of Hold Time (assuming CASW == 1). Thus very fast systems or very noisy systems may want to extend the Address Hold Time after $\overline{\text{CAS}}$ writes to 1.5 clocks. This can be done indirectly by changing the DramWrBTA field. (Address Hold Time on DRAM reads is always at least 1.5 clocks). The default at reset is the value '0' which decodes to 0 clocks of Write Cycle BTA. Field encodings are listed in Table 10.24.

Value	Action
'1'	1 clock
'0'	0 clocks (default)

Table 10.24 DRAM Write Cycle Bus Turn-Around ('DramWrBTA') Field Encoding.

Burst Acknowledge Placement ('DramBurstAck') Field (bit 4:0)

On 4-word burst reads, the acknowledge back to the CPU core needs to be placed so that the CPU pipeline can restart optimally. Acknowledge should be placed 3 clock cycles before the last datum arrives. Wait states via $\overline{\text{SysWait}}$ delay the next de-assertion of $\overline{\text{CAS}}$, however BurstAck may have already asserted. Thus on burst memory cycles where $\overline{\text{SysWait}}$ may potentially be asserted, BurstAck must be programmed to '31' such that it asserts with the last Datum. Field encodings are listed in Table 10.25.

As a reference point, it is from the clock cycle that $\overline{\text{CAS}}$ first asserts and theDramBurstAck internal counter begins counting. The DRAM controller uses the page hit case with no extra $\overline{\text{CAS}}$ Precharge as its minimum value and in the case of page misses and/or in the case of extra $\overline{\text{CAS}}$ Precharge settings, automatically delays the DramBurstAck internal counter.

Note: In the Debug mode of the R36100, if the DebugFCMN pin is asserted, BAck is always automatically asserted with the last Datum.

Value	Action
'31'	Acknowledge with last Datum (use if SysWait is to be asserted).
'30'... '0'	Acknowledge from 30 to 0 clocks (referenced to CAS first asserting). Default is 4.

Table 10.25 Burst Read Acknowledge ('DramBurstAck') Encoding.

Timing Diagrams

The timing diagrams for the R36100 DRAM Controller are divided into the following six sections:

- basic reads
- basic writes
- interleaved reads
- interleaved writes
- transfer mode reads and writes
- refreshes

In the Basic Reads and Basic Writes sections, ordinary 8/16/32-bit DRAM accesses are discussed. Concepts including single versus multiple datum accesses, and many of the option fields, including RASP, RASAd-drHold, AddrSetup, and CASW are shown.

The Interleaved Reads and Interleaved Writes sections discuss the timing for connecting two banks of 32-bit DRAM such that for burst read accesses, both even and odd words are accessed simultaneously and thus improve the performance of the DRAM system.

The Transfer Mode timing section will show how the R36100 implements the Video DRAM Read Output Enable strobe to assert a little earlier than standard DRAM.

In the Refresh timing section, the R36100 is shown to implement a staggered refresh cycle using the $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ protocol of standard DRAMs.

Standard DRAM Chip Summary

The R36100 DRAM Controller uses standard page mode DRAM chips. These DRAM chips multiplex their address pins, support page-mode, and use $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh. Thus on the first part of a DRAM access, upper half of the address is strobed in with a Row Address Strobe ($\overline{\text{RAS}}$) and on the last part of the access, the lower half of the address is strobed in with a Column Address Strobe ($\overline{\text{CAS}}$).

Because page mode DRAMs have an internal array such that an entire row (page) of memory data cells are selected, once that row (page) is turned on, subsequent accesses using $\overline{\text{CAS}}$ can be done much quicker than the first access of any given row.

Thus on burst accesses, the R36100 DRAM Controller keeps $\overline{\text{RAS}}$ asserted and toggles $\overline{\text{CAS}}$ to get multiple datum. However, if a new row (page) is accessed, then the DRAM row array must be re-precharged, typically for two clock cycles. Similarly, $\overline{\text{CAS}}$ is toggled in order to precharge the $\overline{\text{CAS}}$ array before accessing a new memory data location.

On writes, DRAM chips have two modes: early writes and regular writes. Because of the timing advantage of early writes, the R36100 DRAM Controller uses early writes where data is strobed into the DRAM chip with the assertion of $\overline{\text{CAS}}$ instead of with the assertion of the write strobe.

Finally, DRAM devices require that their contents be periodically refreshed. One method is simply to make sure each DRAM row is accessed periodically; however, DRAM chips also have a special $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh protocol: if $\overline{\text{CAS}}$ is asserted before $\overline{\text{RAS}}$, the chip internally executes a refresh access and bumps up an internal row address counter. The R36100 DRAM Controller uses the $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh protocol.

Basic New Page DRAM Read

In Figure 10.10 on page 23, a basic new page DRAM read transaction is shown. The transaction is initiated like other transactions with the assertion of SysALE and SysBurstFrame . Along with the assertion of SysALE , the $\text{SysAddr}()$ bus drives the row address (the upper half of the addresses that the DRAM chips are expecting).

Unlike the Memory Controller, the DRAM Controller has many signals that assert and/or de-assert using the falling edge of SysClk in order to fully optimize the timing for DRAM systems. Thus, 1/2 clock cycle after SysALE asserts, one of the four $\overline{\text{DramRAS}}(3:0)$ strobes will assert depending on which of the four banks is selected. This gives the DRAM chips minimal address setup time to the $\overline{\text{RAS}}$ strobe.

One-half clock cycle after $\overline{\text{DramRAS}}$ asserts, the $\text{SysAddr}()$ bus switches, giving 1/2 clock of address hold time, and begins driving the column address (the lower half of the addresses that the DRAM chips are expecting). One-half cycle after $\text{SysAddr}()$ changes to the column address, from one to four of the $\overline{\text{DramCAS}}(3:0)$ strobes will assert depending on if a particular byte is required on the read.

Note that the R36100 may assert all four $\overline{\text{CAS}}$ lines even though it only requires some of the bytes (in such a case, the unneeded bytes are ignored by the R36100 internally). The default $\overline{\text{CAS}}$ assertion gives the column address setup time to the $\overline{\text{CAS}}$ strobe. In a typical read, $\overline{\text{DramCAS}}()$ remains active for 1.5 clocks. On the final clock of the assertion of $\overline{\text{DramCAS}}()$, SysDataRdy is asserted and the data from the DRAM is latched into the CPU on the final SysClk rising edge.

During the time that $\overline{\text{DramRAS}}(0)$ is active, one of the read enable strobes will also be asserted. These read enable strobes, either $\overline{\text{DramRdEnEven}}$ (see Figure 10.10) or $\overline{\text{DramRdEnOdd}}$ (asserted analogously to the 'Even' signal) can be used to select even ($\overline{\text{DramRAS}}(2)$ or $\overline{\text{DramRAS}}(0)$) or odd ($\overline{\text{DramRAS}}(3)$ or $\overline{\text{DramRAS}}(1)$) memory banks, respectively when multiple banks or transceivers are used. The use of $\overline{\text{DramRdEn}}(\text{Even/Odd})$ varies slightly depending on the type of transceivers and interleaving factor and will be explained in later sections of this chapter. This includes "Interleaved Reads," "Interleaved Writes," and "System Examples."

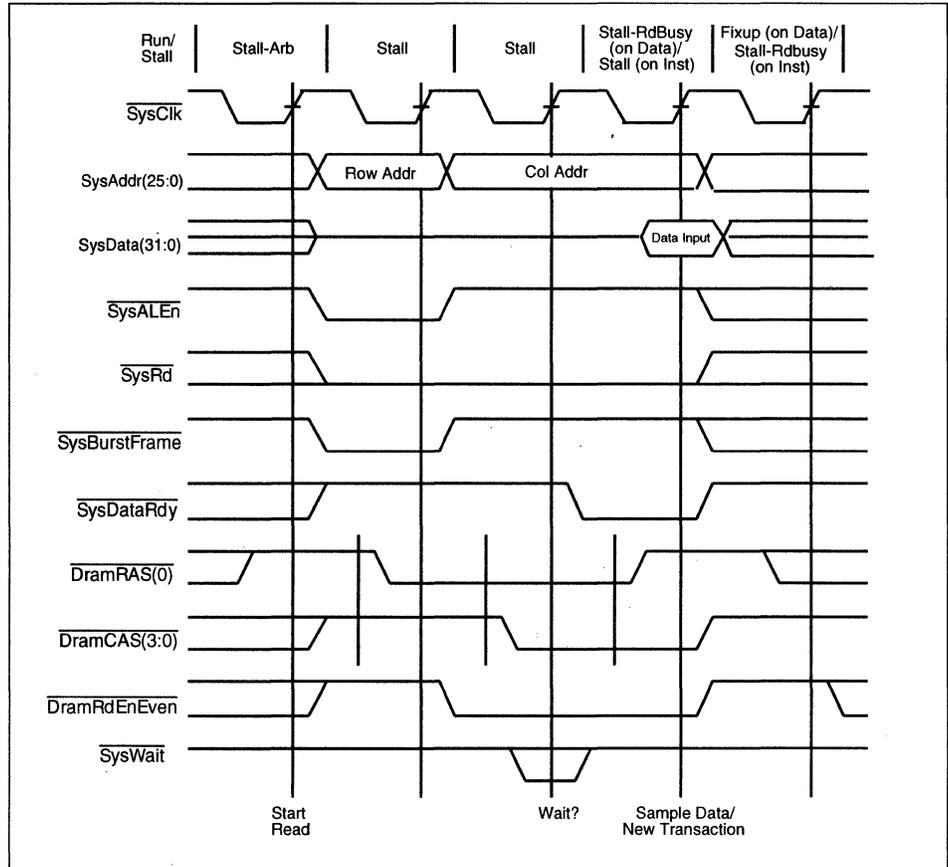


Figure 10.10 Basic DRAM Read

Note: DRAM styles FCT245 and FCT260 have slightly different $\overline{\text{DRAMRdEnEven}}$ behavior than the case shown in Figure 10.10.

$\overline{\text{RAS}}$ Asserted at End of Transfer

The R36100 DRAM Controller contains programmable options to allow $\overline{\text{DramRAS}}(0)$ to be left asserted at the end of a DRAM transaction. Leaving $\overline{\text{RAS}}$ asserted allows a subsequent DRAM transaction to go directly into the $\overline{\text{CAS}}$ stage if the next transaction is to the same row (page) as the previous transaction. Thus by using the Control LSB Register Page Type ('PType') field, the DRAM Controller can keep $\overline{\text{RAS}}$ asserted after burst reads, single word reads, and/or writes. The DRAM Controller accomplishes this by using its internal Page Comparator as described in the Page Comparator Algorithm section earlier in this chapter. Figure 10.11 illustrates this operation.

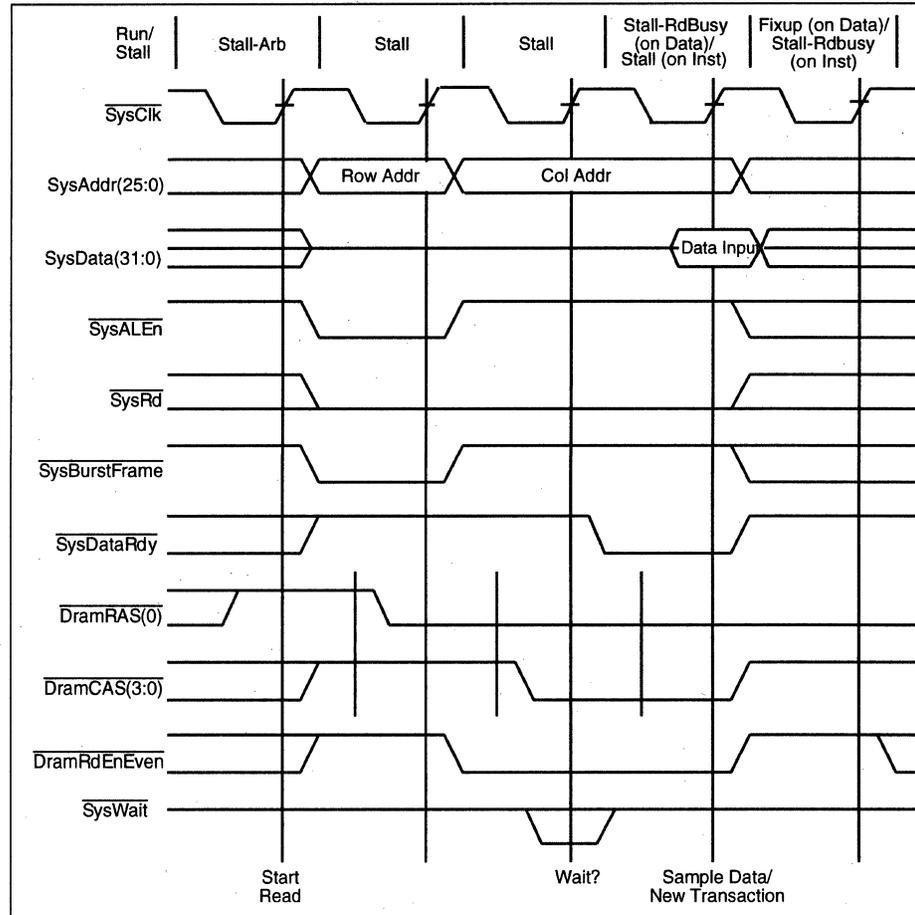


Figure 10.11 RAS asserted at End of Transfer

RAS Asserted at Start of Transfer

If the R36100 DRAM Controller leaves $\overline{\text{DramRAS}}(0)$ asserted at the end of a previous DRAM transaction and the current DRAM transaction is on the same row (page), then $\overline{\text{DramRAS}}(0)$ does not go through a precharge/address strobe stage and is skipped. In this case as shown in Figure 10.12, the $\overline{\text{CAS}}$ address strobe and data access stages happen immediately at the start of the transaction. Note that intervening non-DRAM accesses do not affect the page comparator.

RAS Asserted Throughout Transfer

If the R36100 DRAM Controller leaves $\overline{\text{DramRAS}}(0)$ asserted at the end of a previous transaction and the current transaction is on the same row (page), then $\overline{\text{DramRAS}}(0)$ does not go through a precharge/address strobe stage and is skipped. In this case, as shown in Figure 10.12, the $\overline{\text{CAS}}$ address strobe and data access stages happen immediately at the start of the transaction. As in the RAS Asserted at End Case discussed above, the R36100 DRAM Controller contains programmable options to allow $\overline{\text{DramRAS}}(0)$ to be left asserted at the end of a DRAM transaction.

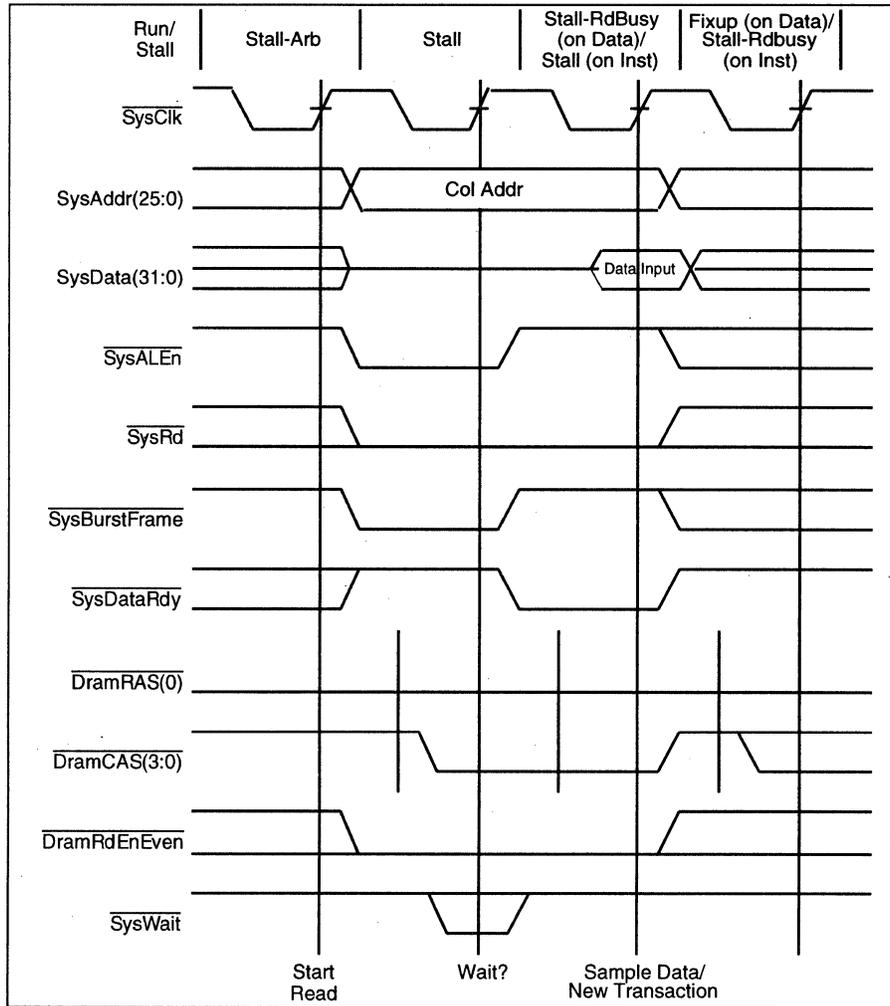


Figure 10.12 RAS asserted at Start of Transfer

RAS Precharge Field

The R36100 DRAM Controller generates an access that either has $\overline{\text{RAS}}$ asserted on a row (page) or has $\overline{\text{RAS}}$ de-asserted. On a subsequent access to a different row (page), the DRAM Controller then verifies that $\overline{\text{RAS}}$ either is de-asserted or has been de-asserted for at least an amount of clocks equal to the Control Register 0 $\overline{\text{RAS}}$ Precharge ('RASP') Field.

Figure 10.13 shows a RASP of 2 clocks where $\overline{\text{RAS}}$ was left asserted on the previous DRAM transaction. To precharge the DRAM chips, $\overline{\text{DramRAS}}(0)$ must first de-assert for 2 clocks. Then $\overline{\text{DRAMRAS}}(0)$ asserts after the 2 clocks, and the transaction continues.

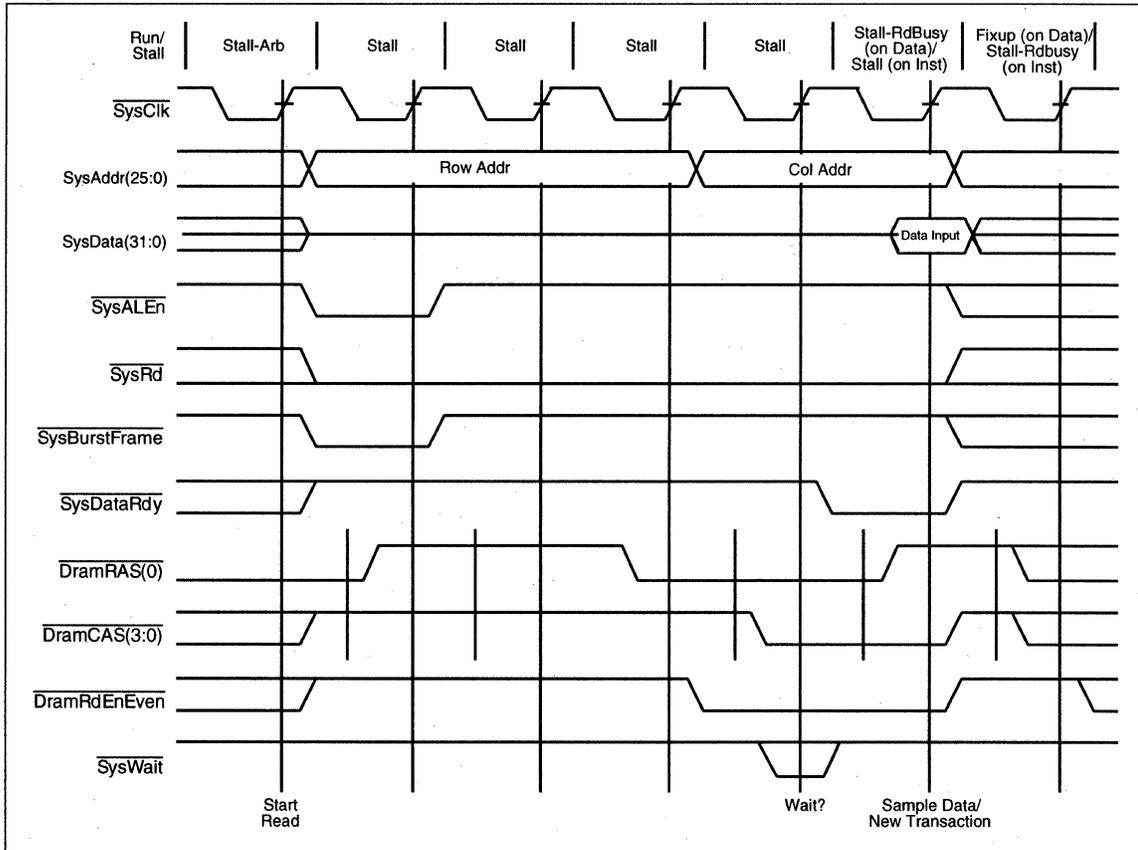


Figure 10.13 RAS Precharge at start of Transfer

RAS Address Hold Field

The RASAddrHold setting can provide extra row address hold time by extending the number of clocks that the DRAM Address Multiplexer delays before switching between row and column addresses. Figure 10.14 shows a DRAM read where RASAddrHold has been set for 1.5 clocks instead of the default of 0.5 clocks in fast or noisy systems, as shown in Figure 10.14.

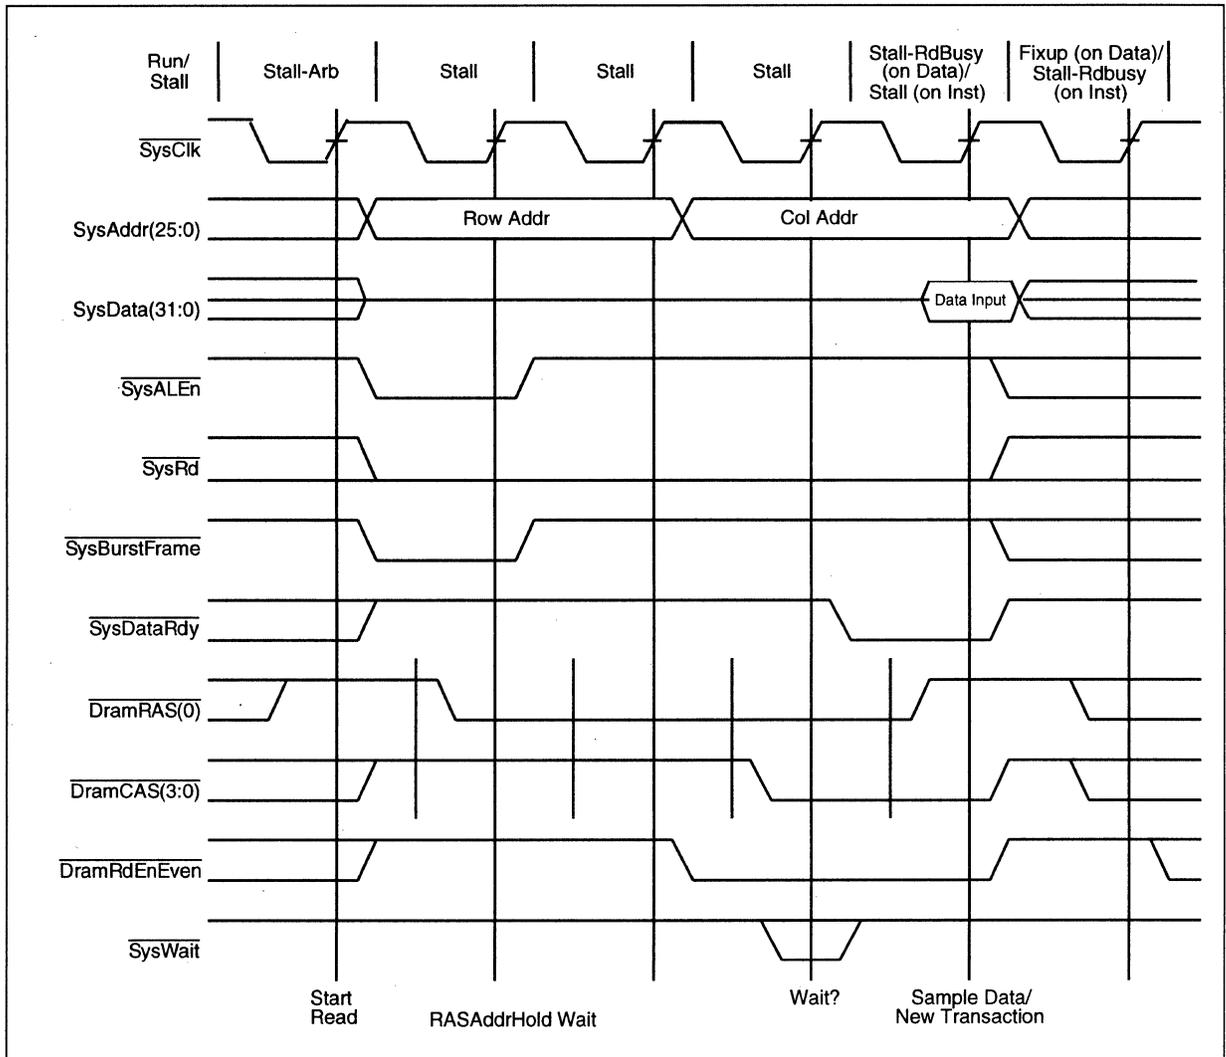


Figure 10.14 Extended Row Address Hold

Address Setup Field

Whenever the R36100 DRAM Controller generates an access where $\overline{\text{RAS}}$ is de-asserted and a new $\overline{\text{RAS}}$ is generated (new page case) or a new $\overline{\text{CAS}}$ is generated, the DRAM Controller is responsible for making sure that $\overline{\text{RAS}}$ or $\overline{\text{CAS}}$ is de-asserted for at least the DRAM Control Register Addr Setup Field amount of clocks after the row or column address is valid.

Figure 10.15 shows a case with AddrSetup of 1.5 clocks between multiple datum on a read. Although this field primarily controls the address setup time of $\overline{\text{CAS}}$ relative to the column address being valid, this field also allows control over the precharge time before $\overline{\text{DramCAS}}(0)$ asserts. To match the column address setup time characteristic, the R36100 DRAM Controller also applies this field to the $\overline{\text{DramRAS}}(0)$ signal relative to the row address on cases where $\overline{\text{DramRAS}}(0)$ was left de-asserted from a previous transaction.

Although not pictured, the AddrSetup field also applies to $\overline{\text{CAS}}$ in the case where $\overline{\text{RAS}}$ is left asserted and then a subsequent same page access occurs. The RASPrecharge field takes care of the case where $\overline{\text{RAS}}$ is asserted and then a subsequent different page access occurs.

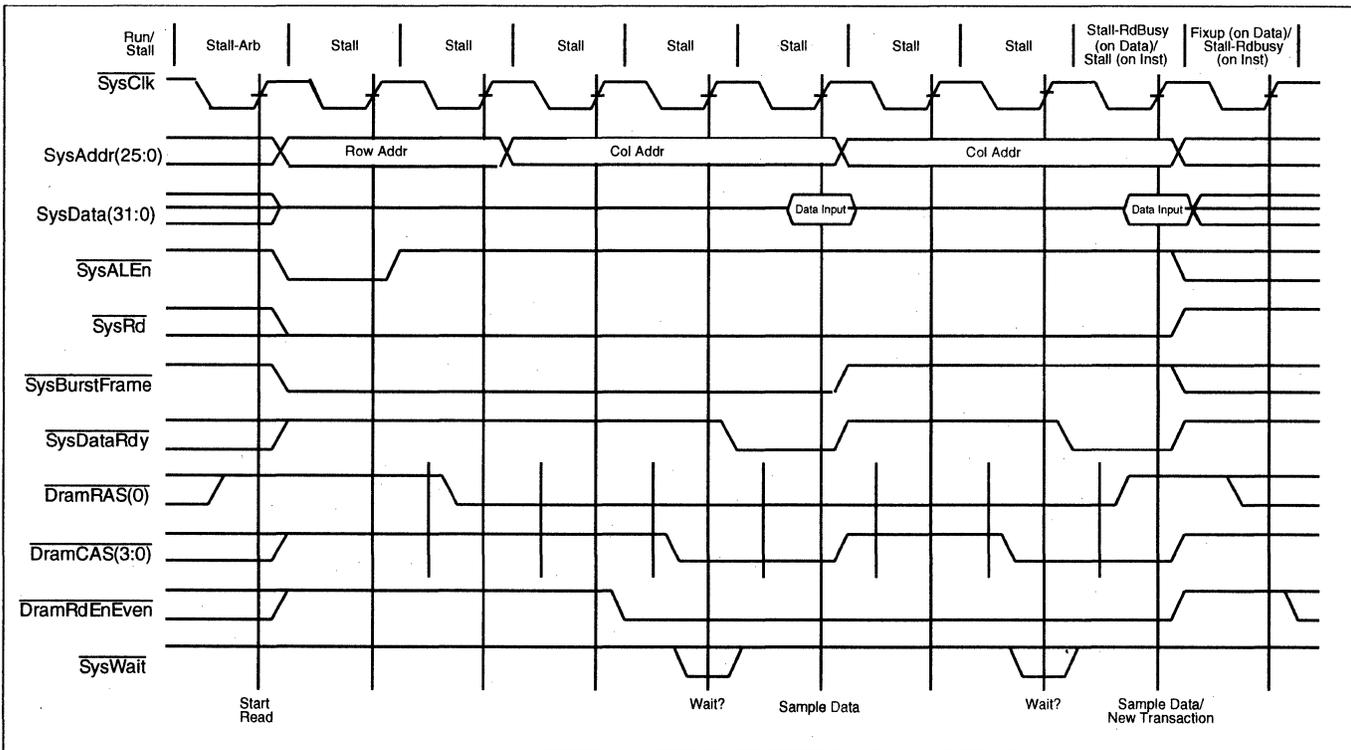


Figure 10.15 Extended Address Set-up

CAS Width Field

The R36100 DRAM Controller can support slower DRAM speeds by increasing the $\overline{\text{CAS}}$ pulse width. This option is programmable using the DRAM Control Register $\overline{\text{CAS}}$ Width ('CASW') field. Figure 10.16 shows the case where CASW has been set to 2.5 clocks instead of the default 1.5 clocks.

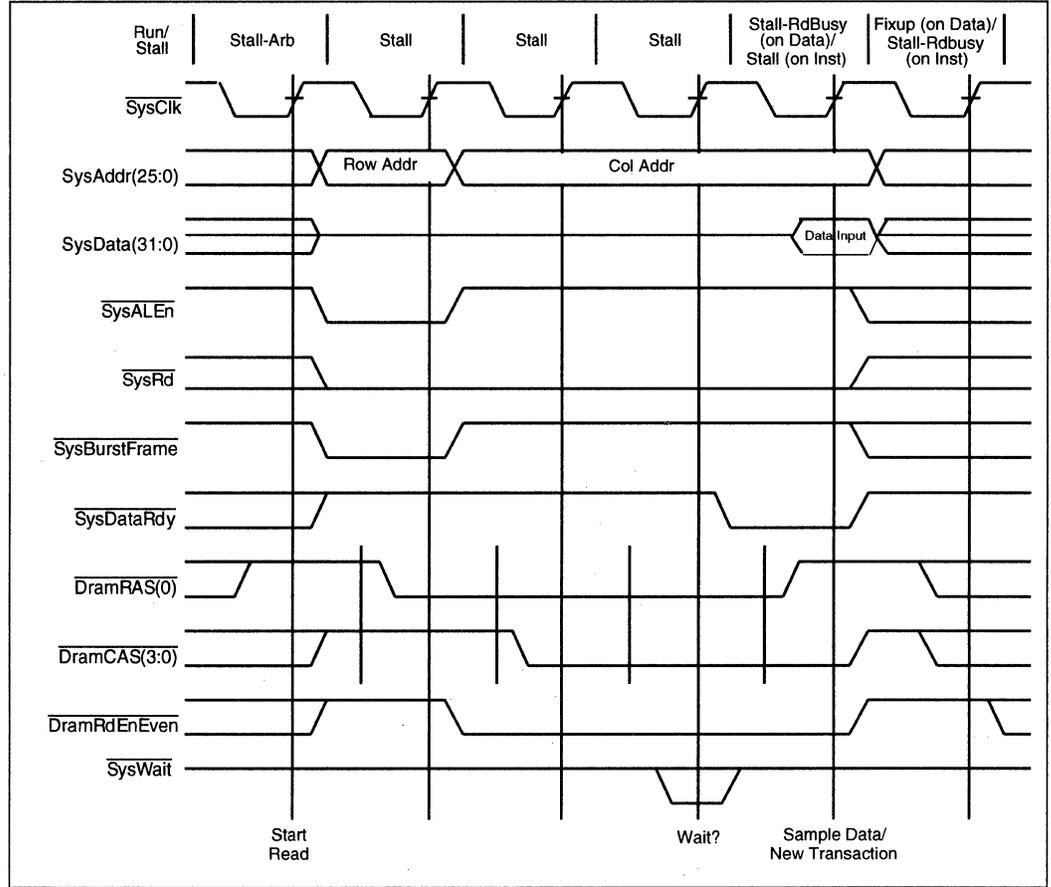


Figure 10.16 Extended CAS Width

Multiple Datum Reads

The R36100 DRAM Controller groups mini-bursts (word and tri-byte accesses on a 16-bit wide port) and non-interleaved bursts (4-word cache refill) the same way. Mini-bursts and burst reads require multiple datum. As shown in Figure 10.17, second and subsequent datum are first preceded with $\overline{\text{DramCAS}}(0)$ de-asserting for 1/2 clock (default) and the DRAM mux'ed SysAddr counting up towards the next column address. (On 16-bit ports A1 is the LSB; on 32-bit ports A2 is the LSB). In each case, the final datum is denoted by $\overline{\text{SysBurstFrame}}$ de-asserting. Although not shown, AddrSetup and CASW fields apply to each $\overline{\text{CAS}}$ Datum.

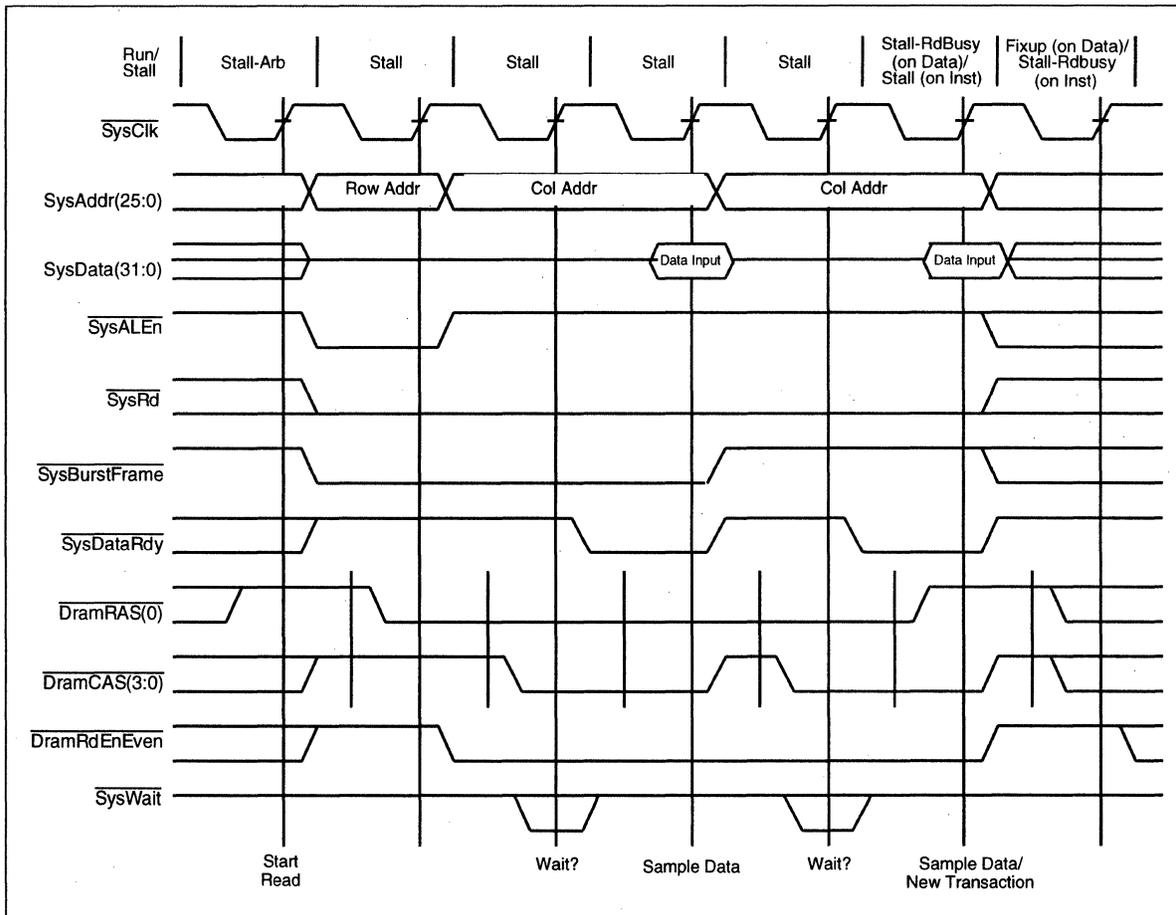


Figure 10.17 Multiple Datum read

Basic DRAM Write

Figure 10.18 shows a basic DRAM write transaction on a precharged (e.g., after reset or a refresh) new page (row). The transaction is initiated like other transactions with the assertion of $\overline{\text{SysALE}}$ and $\overline{\text{SysBurstFrame}}$. Along with the assertion of $\overline{\text{SysALE}}$, the $\overline{\text{SysAddr()}}$ bus drives the row address (the upper half of the addresses that the DRAM chips are expecting).

Unlike the Memory Controller, the DRAM Controller has many signals that asserted and/or de-assert using the falling edge of $\overline{\text{SysClk}}$ in order to fully optimize the timing for DRAM systems. Thus 1/2 clock cycle after $\overline{\text{SysALE}}$ asserts, one of the four $\overline{\text{DramRAS(3:0)}}$ strobes will assert (depending on which one of the four banks is selected). This gives the DRAMs address setup time to the $\overline{\text{RAS}}$ strobe. 1/2 clock cycle after $\overline{\text{DramRAS}}$ asserts, the $\overline{\text{SysAddr()}}$ bus switches and begins driving the column address (the lower half of the addresses that the DRAM chips are expecting).

In addition, the $\overline{\text{SysData()}}$ bus begins driving the appropriate data. One-half cycle after $\overline{\text{SysAddr()}}$ changes, from one to four of the $\overline{\text{DramCAS(3:0)}}$ strobes will assert, depending on if a particular byte is required on the write. The default $\overline{\text{CAS}}$ assertion gives the column address setup time and data setup time to the $\overline{\text{CAS}}$ strobe. The DRAM Controller uses the early write mode of page mode DRAMs where the data is latched by the DRAM chips on the asserting edge of $\overline{\text{CAS}}$ instead of the de-asserting edge. Thus $\overline{\text{SysDataRdy}}$ also asserts a clock early, to indicate to external resources, such as a logic analyzer, that data is valid.

The early write mode allows address pipelining if another non-DRAM access is waiting to use the system bus. Because of the early write mode and address pipelining, the data for the write may disappear on the final clock of the write, if the Write Bus Turn Around is programmed to be '0', because another non-DRAM transaction may have already started. To prevent address pipelining on systems that require additional data hold time (either very high frequency systems or very noisy systems), the Write Bus Turn Around can be programmed to be '1'.

During the time that $\overline{\text{DramCAS()}}$ is active, one of the write enable strobes will also be asserted. These write enable strobes, either $\overline{\text{DramWrEnEven}}$ or $\overline{\text{DramWrEnOdd}}$ can be used to select even ($\overline{\text{DramRAS(2)}}$ or $\overline{\text{DramRAS(0)}}$) or odd ($\overline{\text{DramRAS(3)}}$ or $\overline{\text{DramRAS(1)}}$) memory banks, respectively when multiple banks or transceivers are used. The use of $\overline{\text{DramWrEn(Even:Odd)}}$ and $\overline{\text{DramWrEn(Even:Odd)}}$ varies slightly depending on the type of transceivers and interleaving factor and will be further explained in a later section of this chapter, "System Examples."

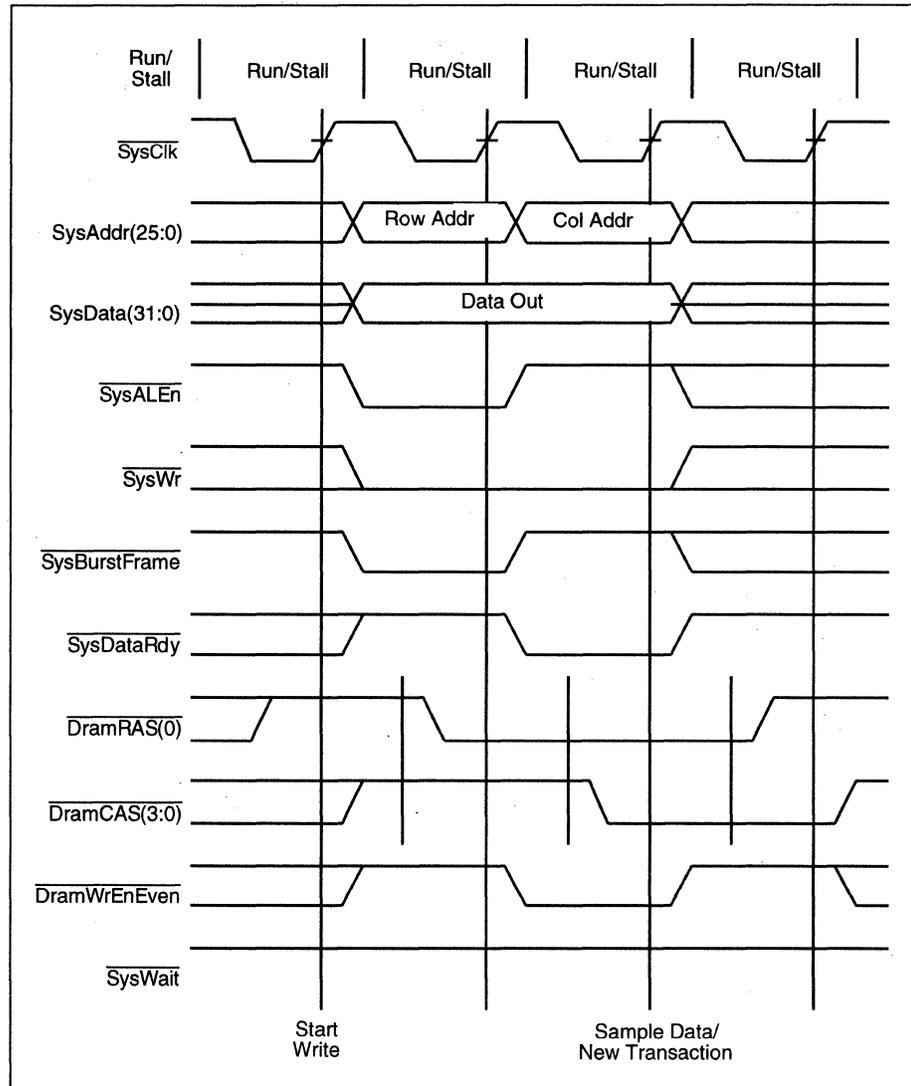


Figure 10.18 Basic DRAM Write

Note: $\overline{\text{DRAMRdEnEven}}$ or $\overline{\text{Odd}}$ (not shown) have slightly different behaviors depending on the DRAM style (FCT245, 260, or 543).

RAS Asserted at Start of Write

If the R36100 DRAM Controller leaves $\overline{\text{DramRAS}}(0)$ asserted at the end of a previous transaction and the current transaction is on the same row (page), then $\overline{\text{DramRAS}}(0)$ does not go through a precharge/address strobe stage and is skipped. In this case, as shown in Figure 10.19, the $\overline{\text{CAS}}$ address strobe and data access stages happen immediately at the start of the transaction.

$\overline{\text{RAS}}$ Asserted at End of Write

The R36100 DRAM Controller contains programmable options to allow $\overline{\text{DramRAS}}_0$ to be left asserted at the end of a DRAM transaction. Leaving $\overline{\text{RAS}}$ asserted allows a subsequent DRAM transaction to go directly into the $\overline{\text{CAS}}$ stage if the next transaction is to the same row (page) as the previous transaction. Thus by using the Control 0 MSB Register Page Type ('PType') field, the DRAM Controller can keep $\overline{\text{RAS}}$ asserted after burst reads, single word reads, and/or writes. The DRAM Controller accomplishes this by using its internal Page Comparator as described in the Page Comparator Algorithm section earlier in this chapter.

RAS Asserted Throughout Write

If the R36100 DRAM Controller leaves $\overline{\text{DramRAS}}(0)$ asserted at the end of a previous transaction and the current transaction is on the same row (page), then $\overline{\text{DramRAS}}(0)$ does not go through a precharge/address strobe stage and is skipped. In this case, as shown in Figure 10.19, the $\overline{\text{CAS}}$ address strobe and data access stages happen immediately at the start of the transaction. As with $\overline{\text{RAS}}$ Asserted at End of Write, the R36100 DRAM Controller contains programmable options to allow $\overline{\text{DramRAS}}(0)$ to be left asserted at the end of a DRAM transaction.

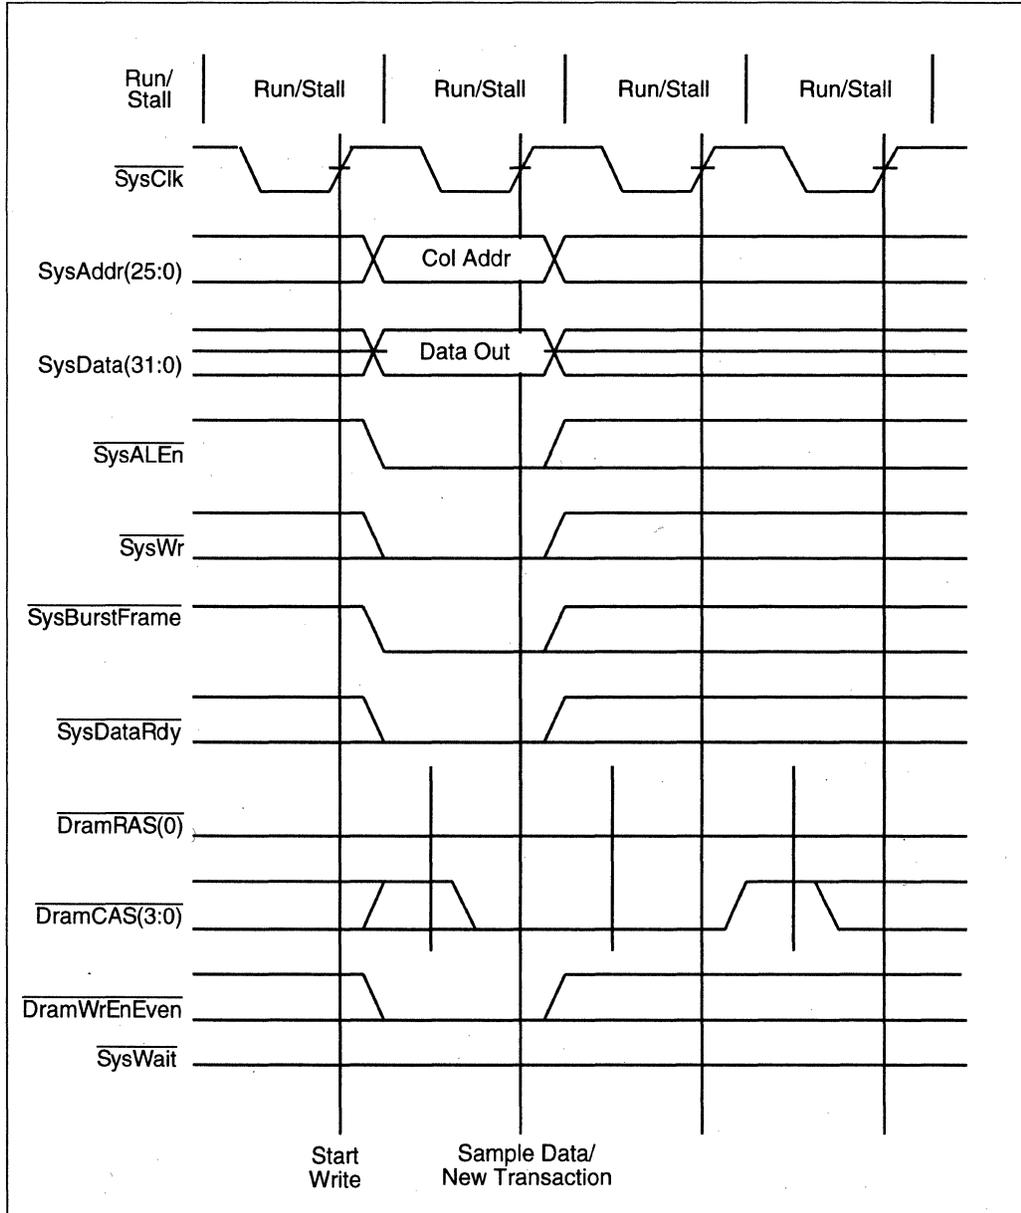


Figure 10.19 $\overline{\text{RAS}}$ Asserted Throughout DRAM Write

Other DRAM timing Controls

Most of the DRAM Control Fields work identically for reads and writes, this includes:

- RAS Precharge Field,
- RAS Address Hold Field,
- Address Setup Field,
- RAS Address Hold Field,
- CASW Field,

For more detail, see the DRAM Read Timing section.

Write Bus Turn-Around

Normally, a subsequent non-DRAM transaction can potentially begin 1 clock before the DRAM has actually completed the write. For example, this DRAM Write Pipelining can occur when a DRAM write is followed by an instruction read from PROM. In some cases where either the system clock frequency is very high or the column address is very noisy, the column address needs additional hold time. By using the Write Bus Turn-Around Field in a DRAM bank's MSB Control Register, the column address is held for an extra clock by delaying any non-DRAM transactions for 1 clock, as shown in Figure 10.20.

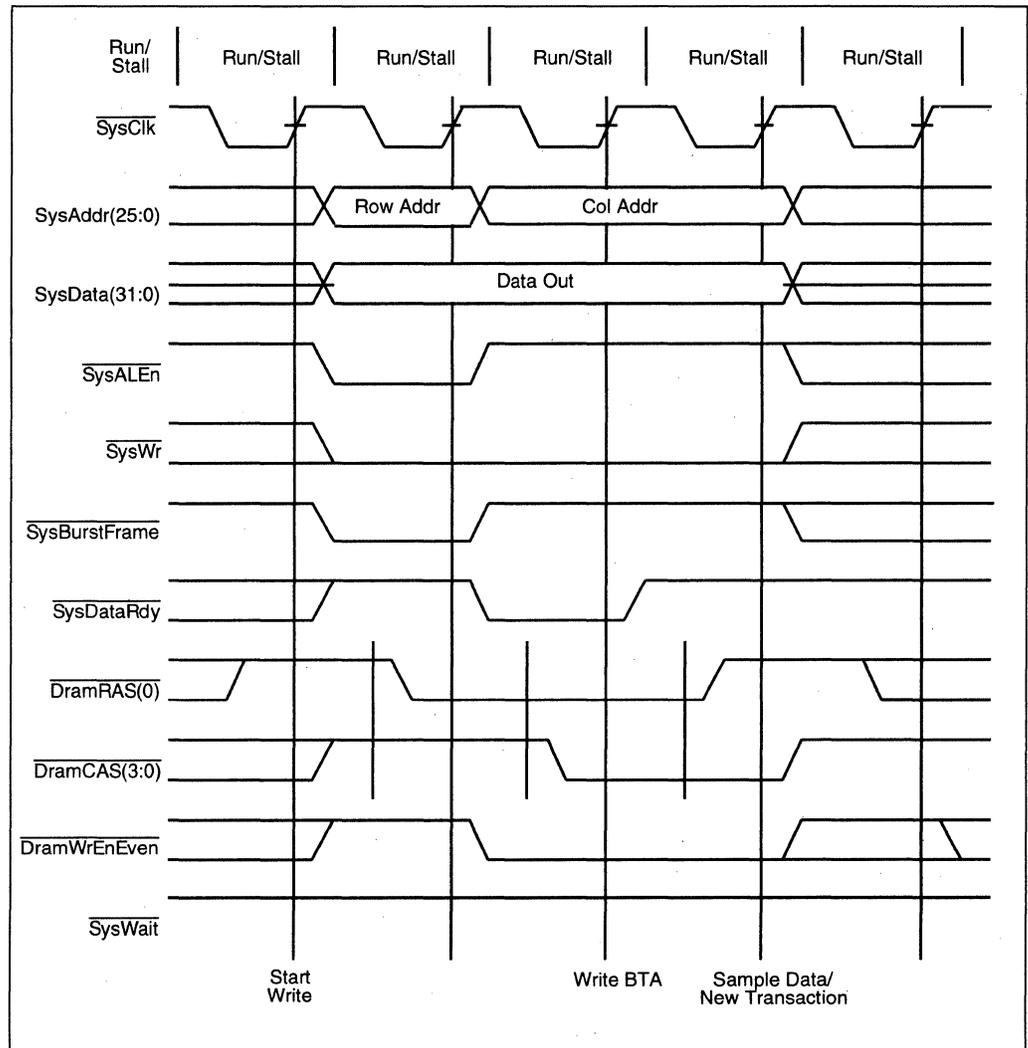


Figure 10.20 Write Bus Turn-around

Two Datum Write Transaction

In cases with a 16-bit bus port width that access more than a halfword (tri-byte, word, or DMA Burst Write) or in cases with a non-interleaved 32-bit bus port width that is a DMA Burst Write, the R36100 DRAM Controller does a two datum or multi-datum write transaction, as shown in Figure 10.21. The second or subsequent data have finished using the DRAM page mode, such that new SysData() is put on the data bus and $\overline{\text{DramCAS}}(0)$ is re-asserted. The control lines, $\overline{\text{DramRdEn}}(\text{Odd/Even})$ and $\overline{\text{DramWrEn}}(\text{Odd/Even})$ for the FCT245-Type transceivers operate slightly differently than for FCT260- or FCT543-Type transceivers as will be explained later this chapter in the sections on “Interleaved Reads,” “Interleaved Writes,” and “System Examples”.

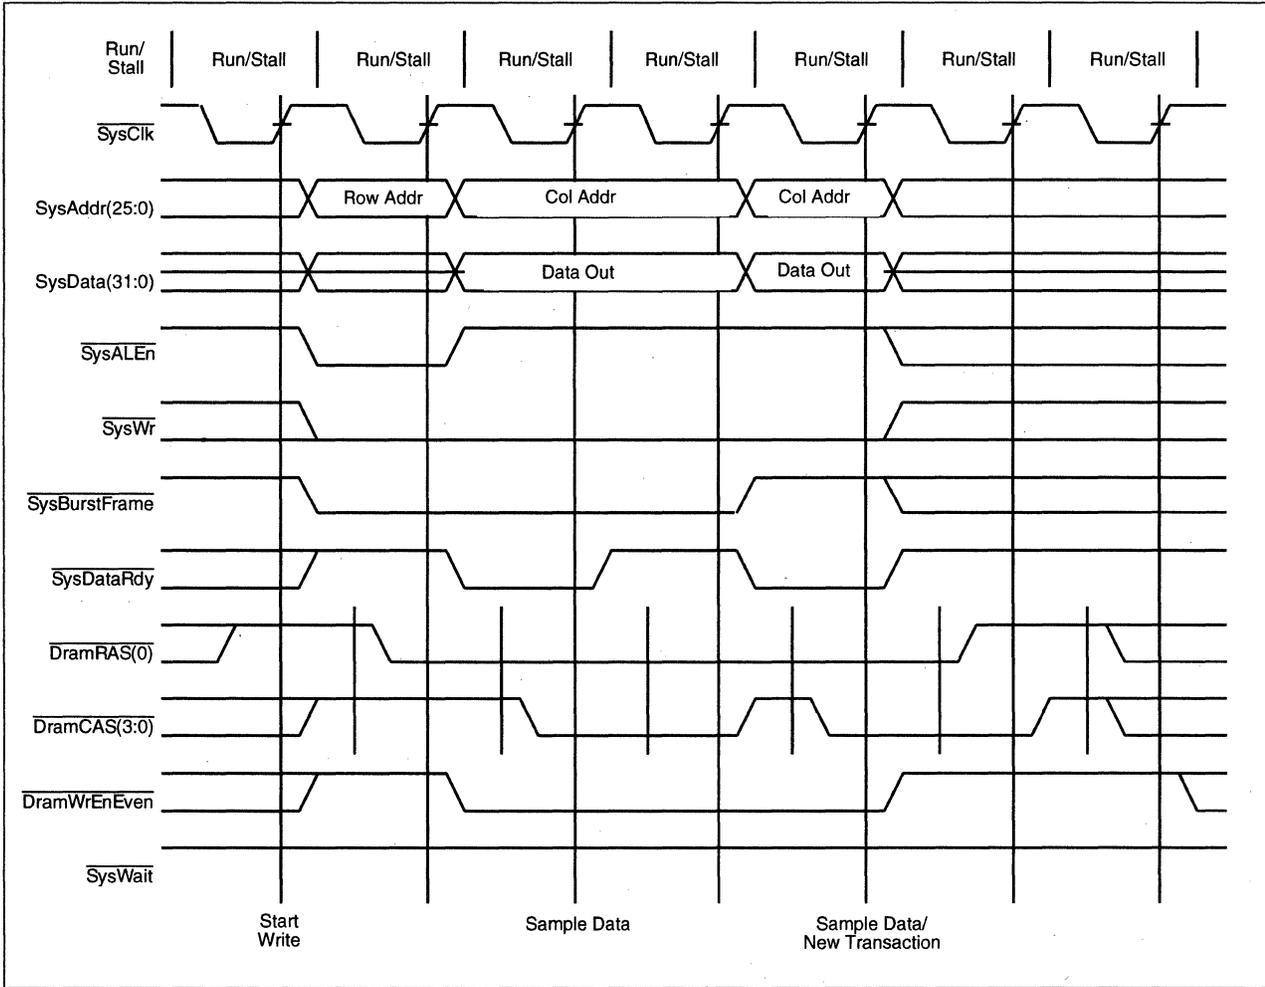


Figure 10.21 Two Datum Write

Interleaved Reads

Interleaved FCT245 Reads

If the DRAM LSB Control Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT245' setting, then the DRAM Controller assumes the timing shown in Figure 10.22 for reads. Because data is not latched by the transceiver, the 2nd, 4th, 6th,... datum must be read with a constant address and CAS assertion. Thus the interleaved FCT245 case saves 1 clock per odd datum over the non-interleaved case. In the Interleaved FCT245 case, the read enables `DramRdEn(Odd,Even)` are used as transceiver enables on both reads and writes. The `DramWrEn(Odd,Even)` signals can be used for the direction.

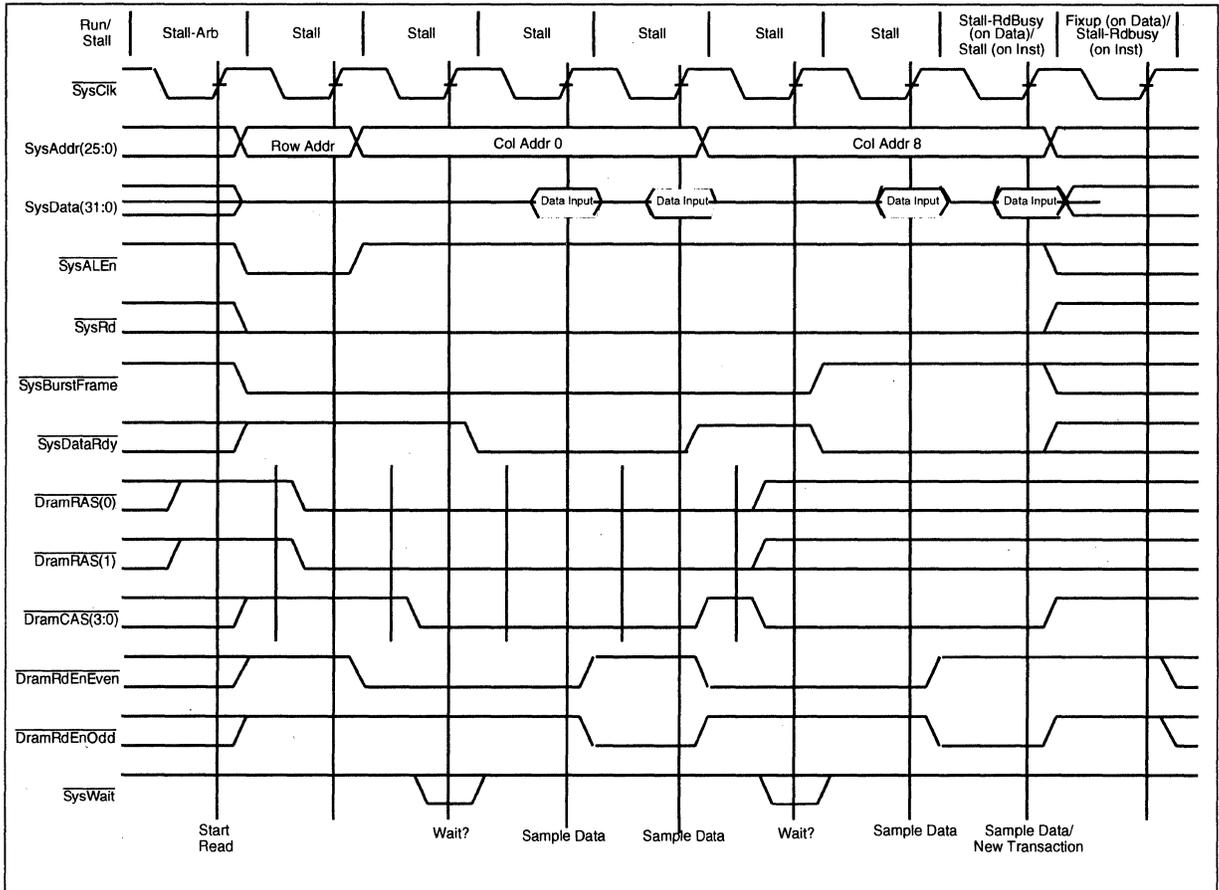


Figure 10.22 Interleaved 'FCT245 type read

Interleaved FCT260 Reads

If the DRAM LSB Control Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT260' setting, then the DRAM Controller assumes the timing shown in Figure 10.23 for reads. It is assumed odd datum are latched by the multiplexer, such that the next address and $\overline{\text{CAS}}$ lines (for the even datum) can be pipelined to change 1 clock early. Thus the interleaved FCT260 case saves at least 3 clocks for each 4-word burst read. In the Interleaved FCT260 case, the even read enable $\overline{\text{DramRdEnEven}}$ is used to latch the odd datum while the odd read enable $\overline{\text{DramRdEnOdd}}$ is used as the overall read enable for the multiplexer.

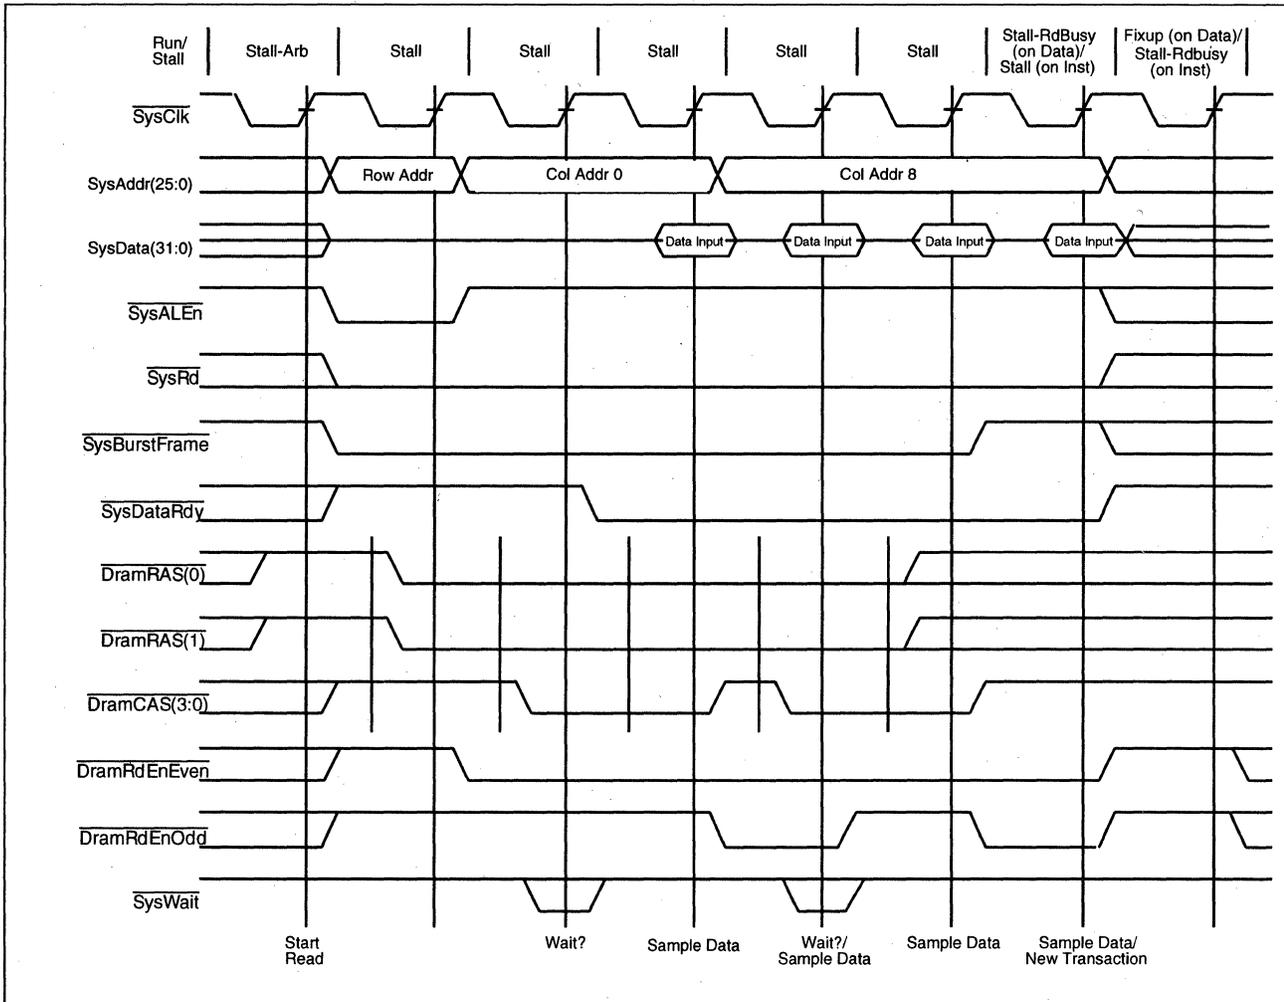


Figure 10.23 Interleaved FCT260 Read

Figure 10.24 shows a single datum access to an interleaved memory system using FCT260-type multiplexers in the data path. Note that the timing of this access is identical to the timing of the first word access of a quad word read.

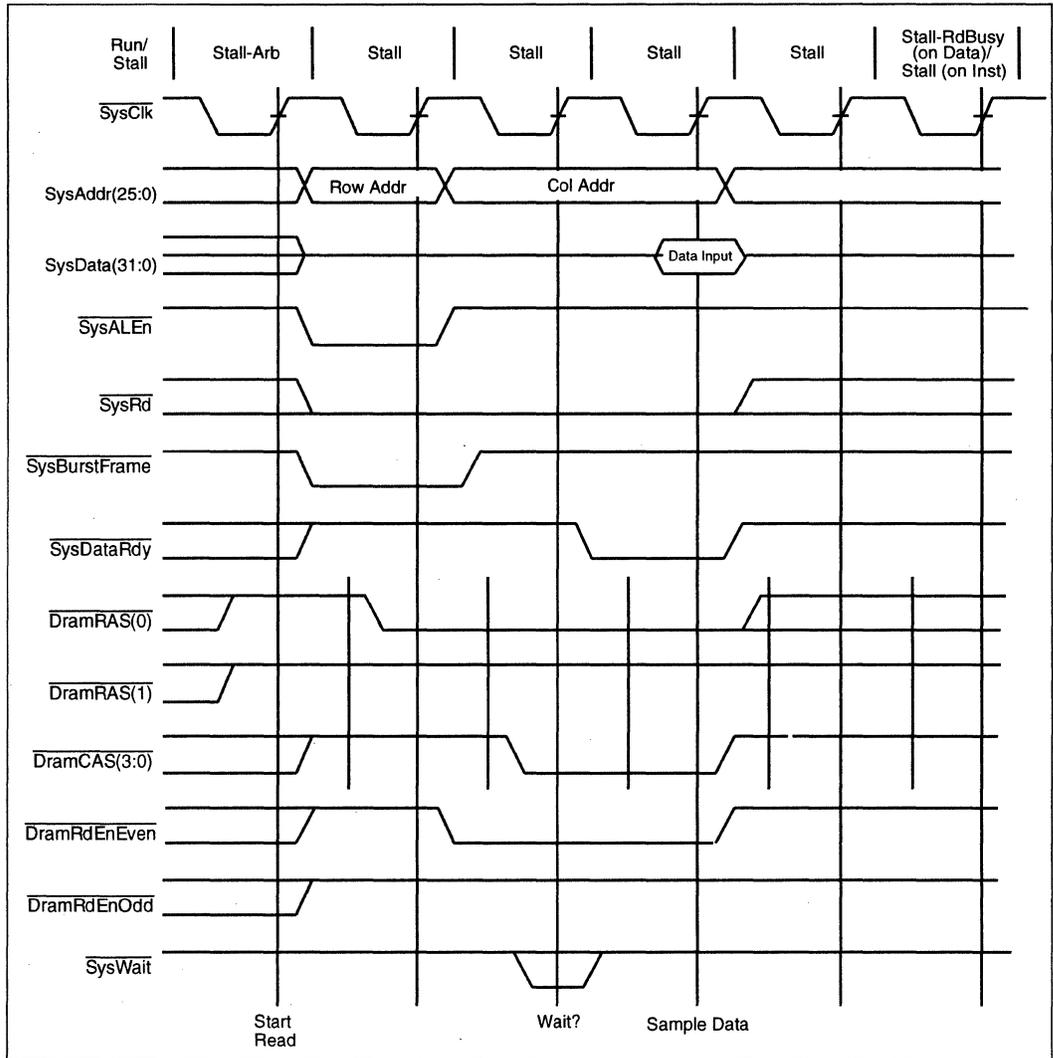


Figure 10.24 Single word access to even bank of FTC260-type system

Figure 10.25 shows the analogous access to the “odd” bank of an interleaved FCT260-type memory system. In this figure, the timing is identical with the timing of the access of the *second* word of a 4-word access; however, the first word is not actually returned to the CPU.

Therefore, there is a performance difference between even and odd single-word accesses, due to a limitation on the number of transceiver control pins available. However, for the following reasons, this should not adversely affect system performance:

- Single word accesses occur for uncached instruction or data fetches. These are typically not used in performance critical parts of the system software.
- Cached instruction misses are always satisfied using 4-word refills, and utilize instruction streaming to resume execution once the critical missing instruction is returned from memory.
- Single word accesses may be used for cached data refills, if the data block refill parameter is set accordingly. However, the use of an interleaved memory in the first place indicates that the burst performance of the memory system is very high, leading to an extremely high probability that 4-word D-cache refill is used.

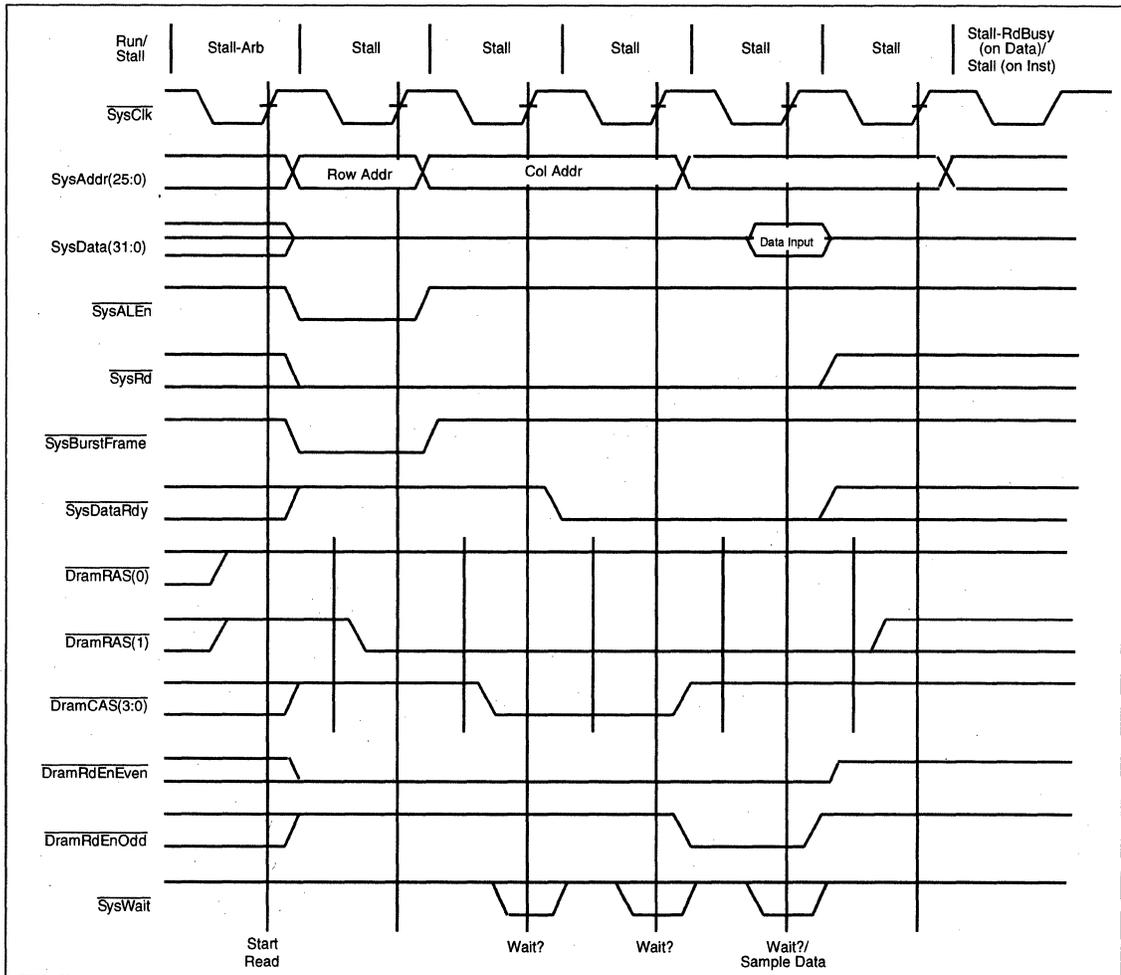


Figure 10.25 Single word access to odd bank of FCT260-type system

Interleaved FCT543 Reads

If the DRAM LSB Control Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT543' setting, then the DRAM Controller assumes the timing shown in Figure 10.26 for reads. It is assumed odd datum are latched by the registered transceiver, such that the next address and $\overline{\text{CAS}}$ lines (for the even datum) can be pipelined to change 1 clock early. Thus the interleaved FCT543 case saves at least 3 clocks for each 4-word burst read. In the Interleaved FCT543 case, the two read enables and two write enables match up with the FCT543 part directly.

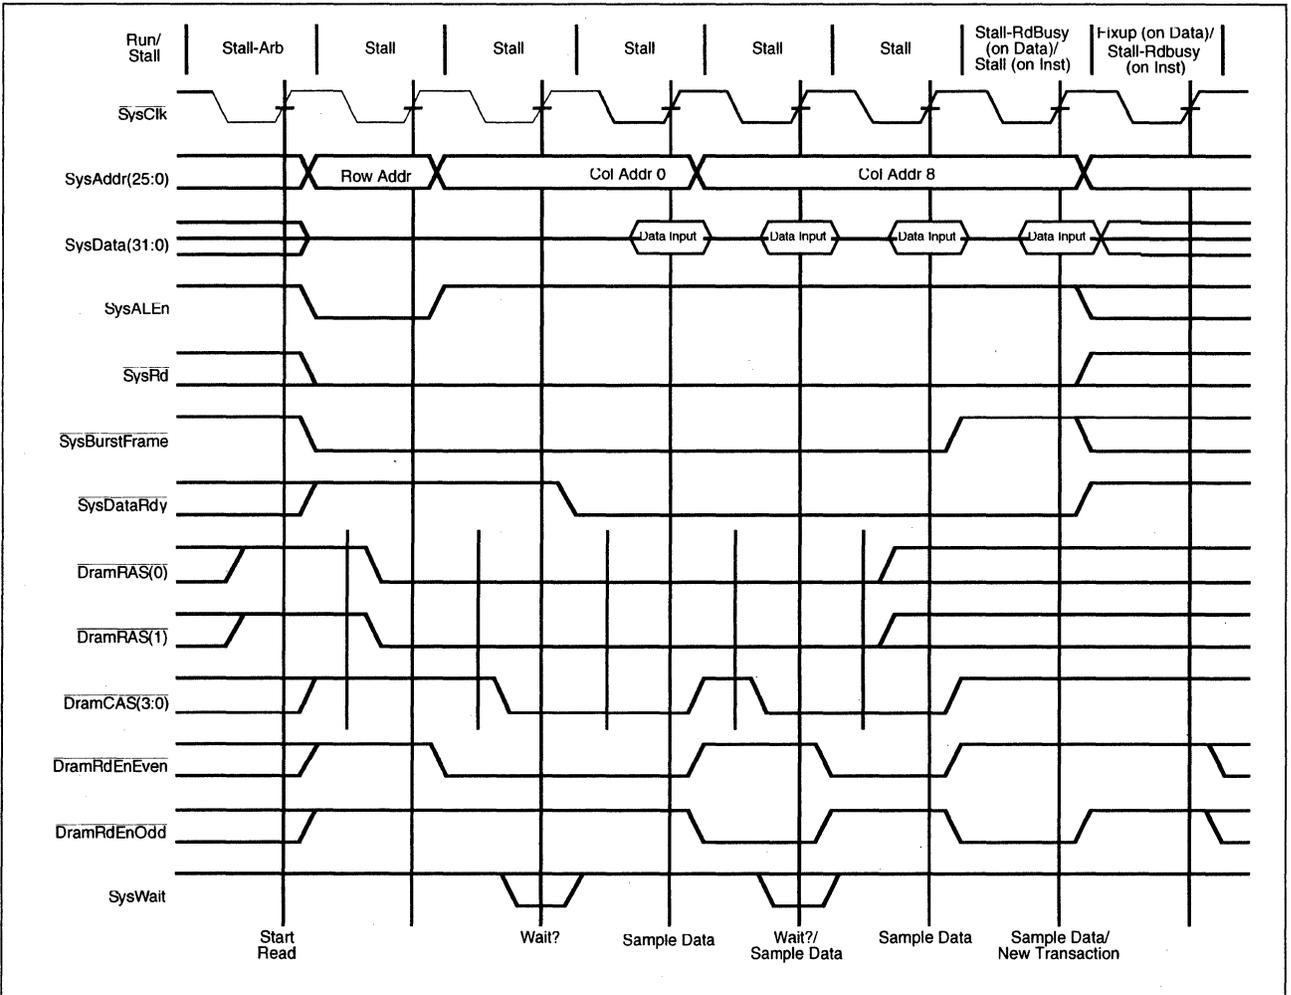


Figure 10.26 Interleaved FCT543 Read

Figure 10.27 shows a single datum access to an interleaved memory system using FCT543-type multiplexers in the data path. Note that the timing of this access is identical with the timing of the first word access of a quad word read.

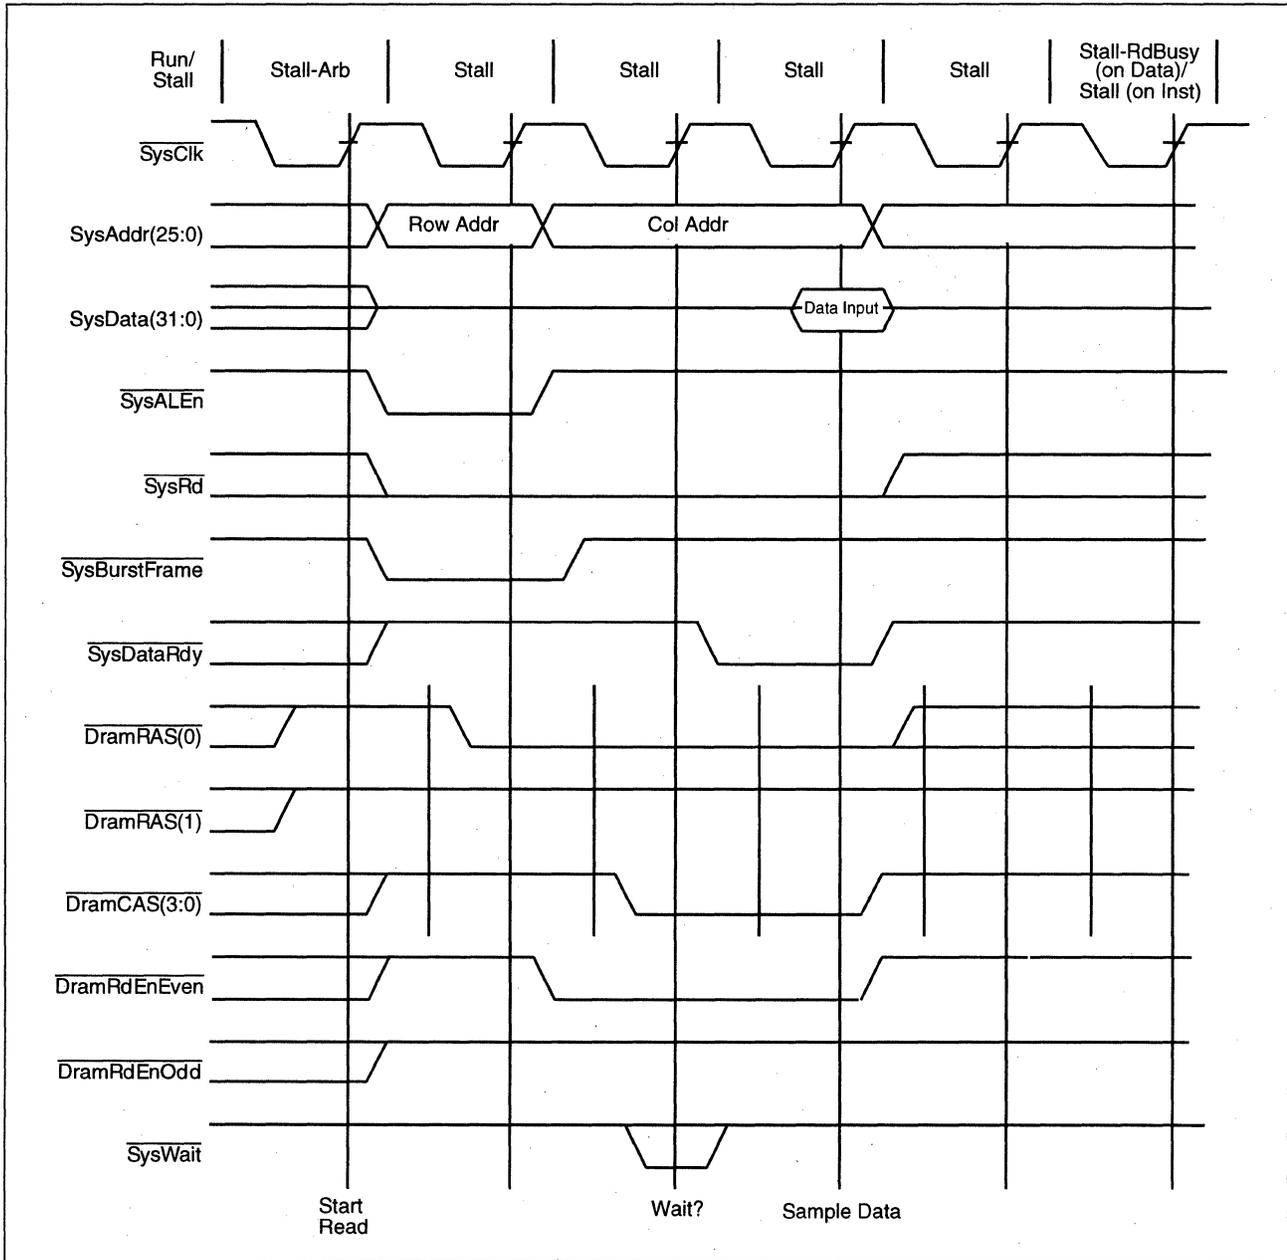


Figure 10.27 Single word access to even bank of FCT543-type system

Figure 10.28 shows the analogous access to the “odd” bank of an interleaved FCT543-type memory system. In this figure, the timing is identical with the timing of the access of the *second* word of a 4-word access; however, the first word is not actually returned to the CPU.

Thus, there is a performance difference between even and odd single word accesses, due to a limitation on the number of transceiver control pins available. However, for the following reasons, this should not adversely affect system performance:

- Single word accesses occur for uncached instruction or data fetches. These are typically not used in performance critical parts of the system software.
- Cached instruction misses are always satisfied using 4-word refills, and utilize instruction streaming to resume execution once the critical missing instruction is returned from memory.
- Single word accesses may be used for cached data refills, if the data block refill parameter is set accordingly. However, the use of an interleaved memory in the first place indicates that the burst performance of the memory system is very high, leading to an extremely high probability that 4-word D-cache refill is used.

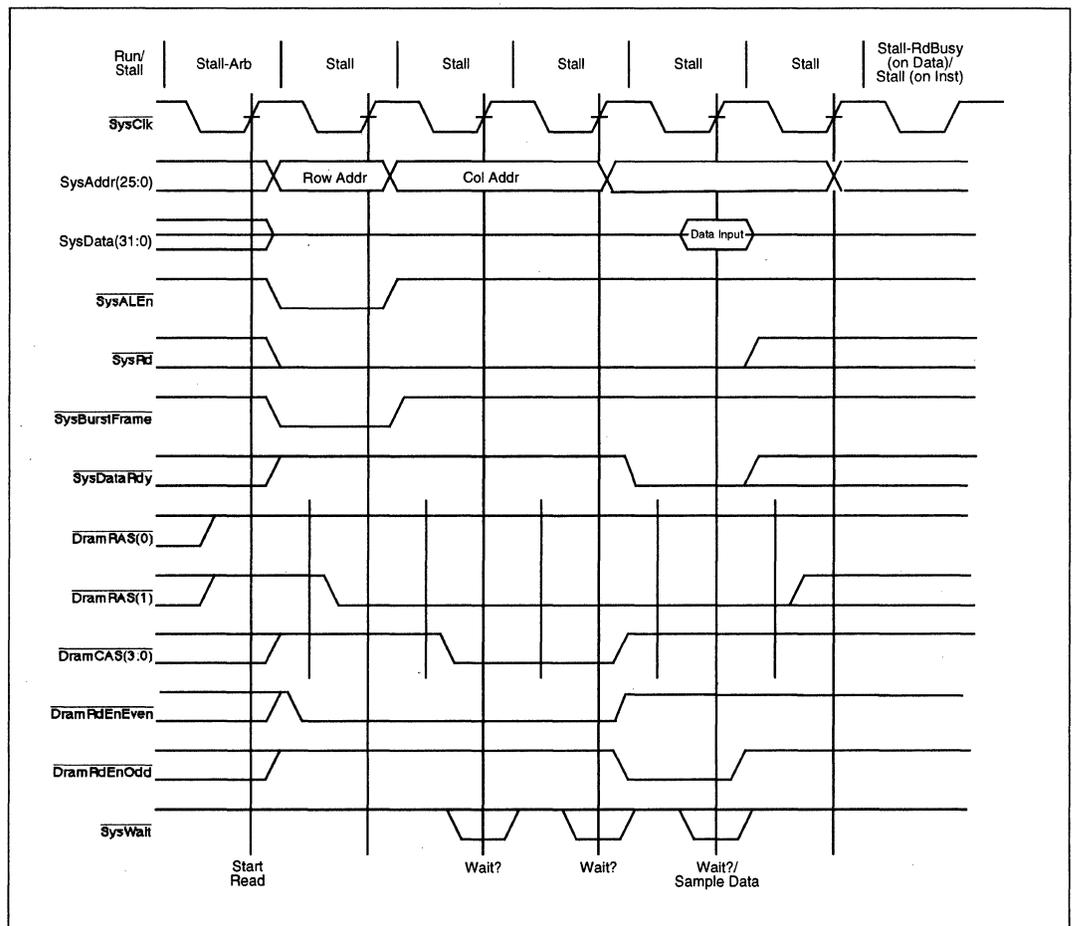


Figure 10.28 Single word access to odd bank of FCT543-type system

Interleaved Writes

Interleaved writes on the R36100 occur one word at a time on the respective bank. The R36100 CPU core is only capable of issuing one write at a time. However, the DMA engines are capable of issuing burst writes. At present, such burst writes are not highly optimized on the R36100 and issue sequentially one after another with separate $\overline{\text{RAS}}$ (as well as $\overline{\text{CAS}}$) strobes, switching between banks. This choice is due to the leading edge of $\overline{\text{CAS}}$ needing to be delayed for early writes on fully optimized bursts, which would cause needless complications for more typical systems.

Single Word Interleaved FCT245 Write

If the DRAM Control MSB Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT245' setting, then the DRAM Controller assumes the timing similar to that shown in the first half of Figure 10.29 on page 45 on writes.

In the Interleaved FCT245 case, the read enables $\overline{\text{DramRdEn(Odd/Even)}}$ are used as transceiver enables on both reads and writes. The $\overline{\text{DramWrEn(Odd/Even)}}$ signals can be used for the direction. The Single Word case is similar to the multi-word case, except that the second assertion of $\overline{\text{CAS}}$ does not occur.

Interleaved FCT245 Writes

If the DRAM Control MSB Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT245' setting, then the DRAM Controller assumes the timing shown in Figure 10.29 on writes.

On interleaved writes, the R36100 DRAM Controller does the writes with 'early writes' and thus if a burst write occurs, separate $\overline{\text{CAS}}$ strobes occur for each datum.

In the Interleaved FCT245 case, the read enables $\overline{\text{DramRdEn(Odd/Even)}}$ are used as transceiver enables on both reads and writes. The $\overline{\text{DramWrEn(Odd/Even)}}$ signals can be used for the direction.

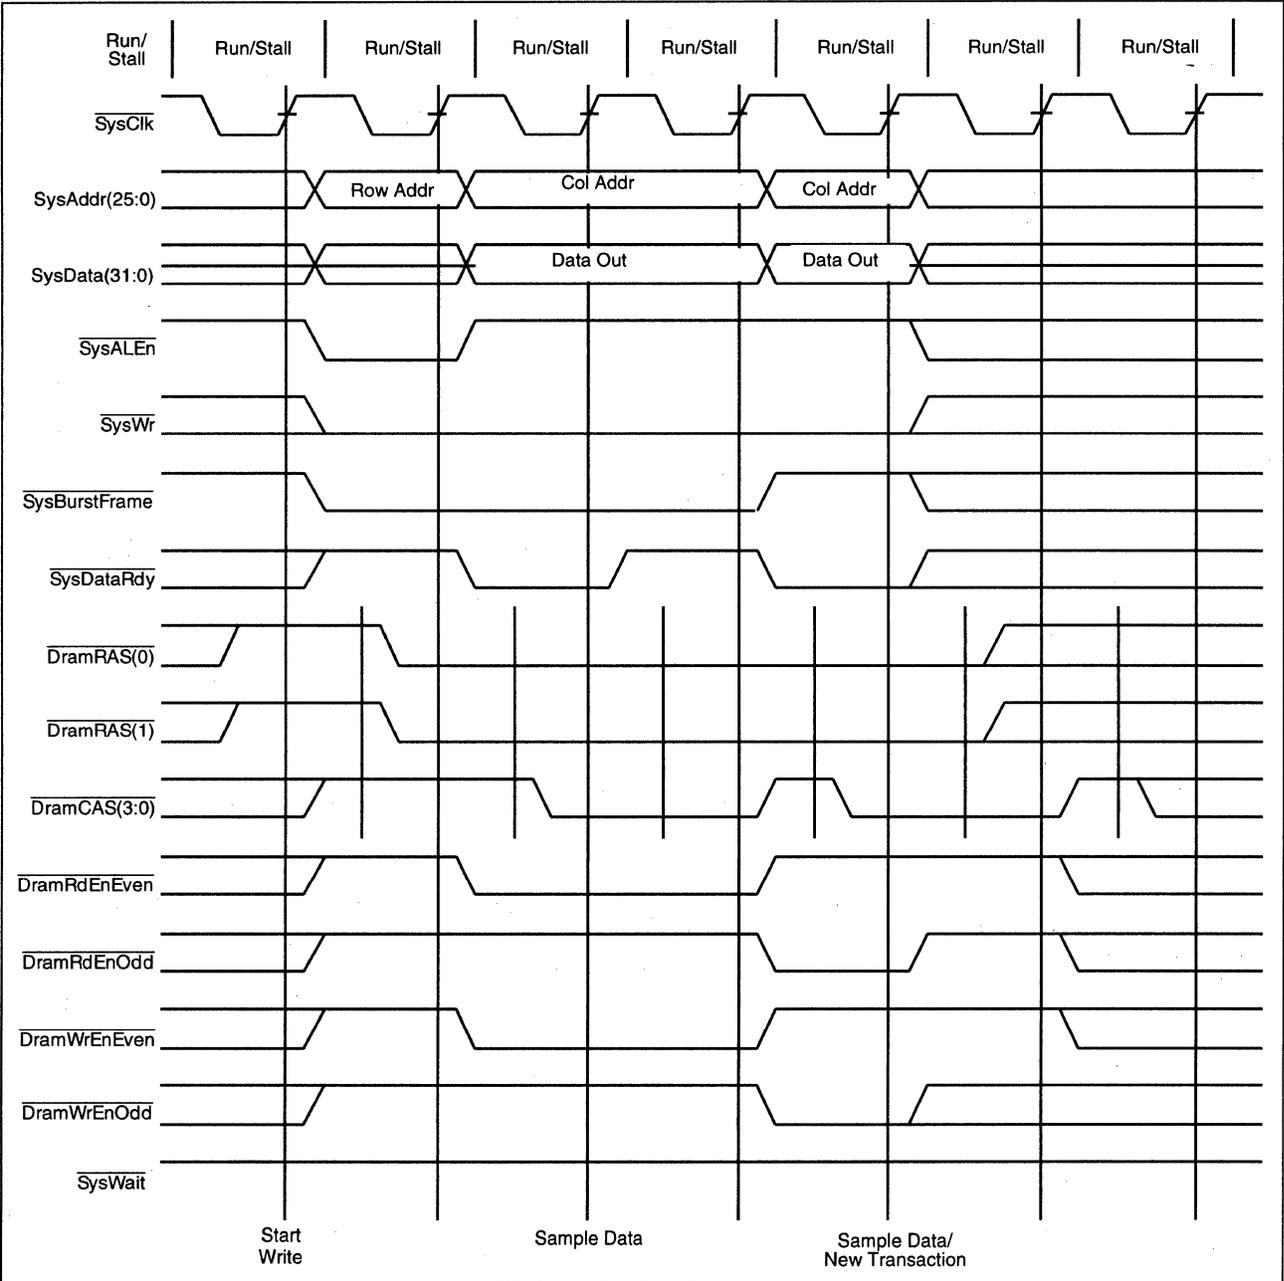


Figure 10.29 Interleaved FCT245-type Writes

Single Word Interleaved FCT260 Write

If the DRAM Control MSB Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT260' setting, then the DRAM Controller assumes the timing shown in on writes.

On interleaved writes, the R36100 DRAM Controller does the writes with 'early writes' and thus if a burst write occurs, separate $\overline{\text{CAS}}$ strobes occur for each datum.

In the Interleaved FCT260 case, the even read enable $\overline{\text{DramRdEnEven}}$ is used to latch the odd datum while the odd read enable $\overline{\text{DramRdEnOdd}}$ is used as the overall read enable for the multiplexer.

The Single Word case is similar to the multi-word case, except that the second assertion of $\overline{\text{CAS}}$ does not occur.

Interleaved FCT260 Writes

If the DRAM Control MSB Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT260' setting, then the DRAM Controller assumes the timing shown in Figure 10.30 on writes.

On interleaved writes, the R36100 DRAM Controller does the writes with 'early writes' and thus if a burst write occurs, separate $\overline{\text{CAS}}$ strobes occur for each datum.

In the Interleaved FCT260 case, the even read enable $\overline{\text{DramRdEnEven}}$ is used to latch the odd datum while the odd read enable $\overline{\text{DramRdEnOdd}}$ is used as the overall read enable for the multiplexer.

Interleaved FCT543 Writes

If the DRAM Control MSB Register Type Field of a pair of chip select channels is programmed to the 'Interleaved FCT543' setting, then the DRAM Controller assumes the timing shown in Figure 10.30 on writes.

On interleaved writes, the R36100 DRAM Controller does the writes with 'early writes' and thus if a burst write occurs, separate $\overline{\text{CAS}}$ strobes occur for each datum.

In the Interleaved FCT543 case, the two read enables and two write enables match up with the FCT543 part directly.

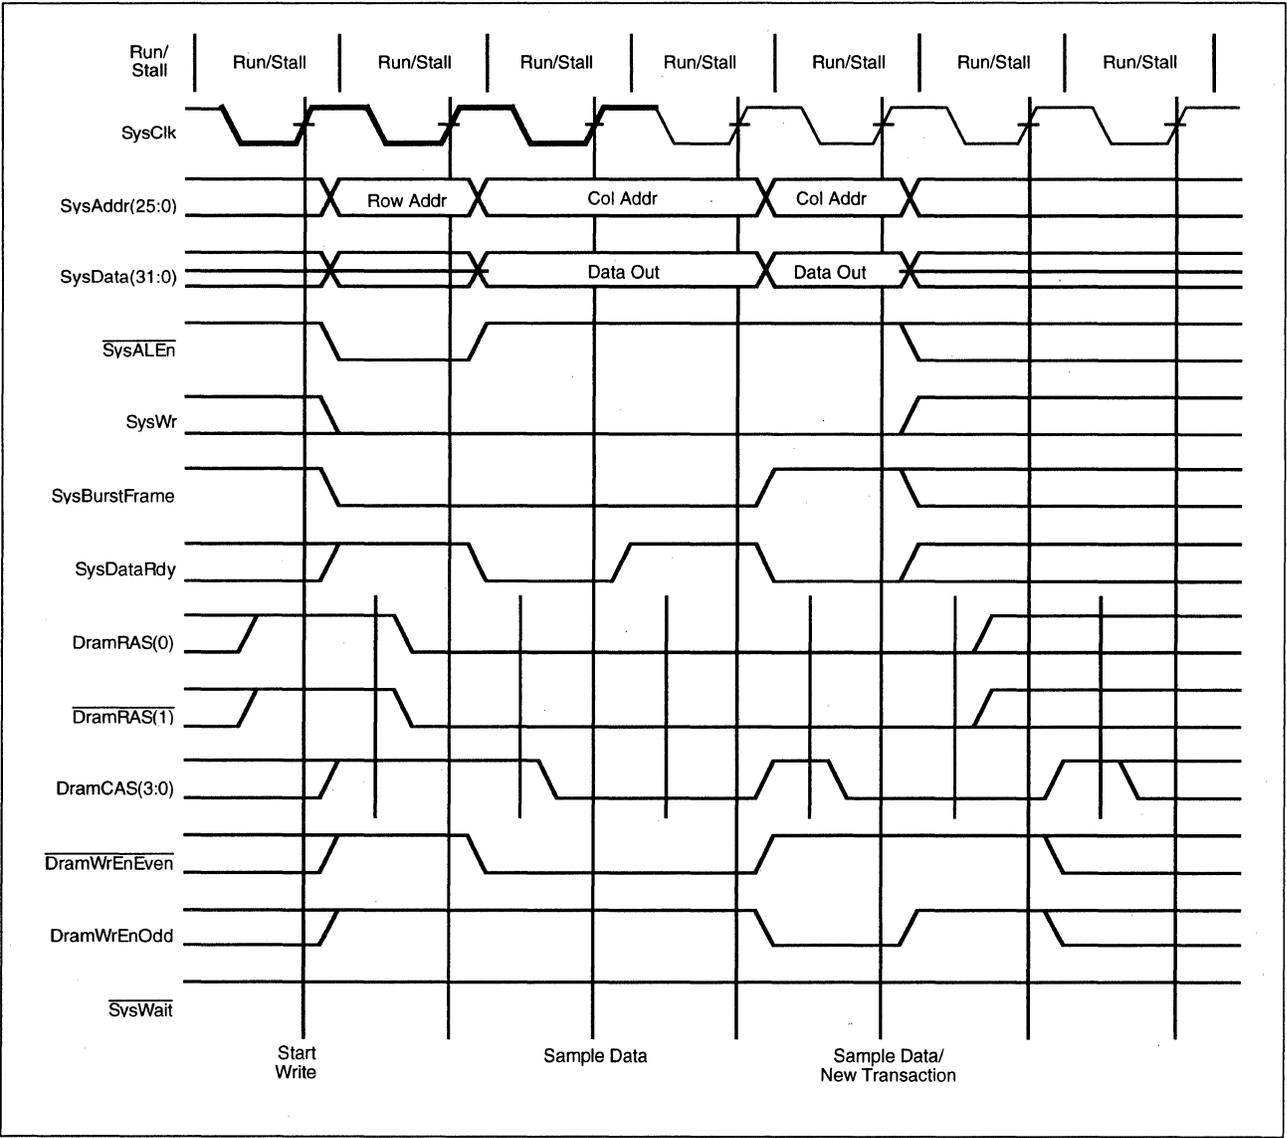


Figure 10.30 Interleaved FCT260, FCT543-type Writes

Refresh

The refresh cycle of DRAM chips is supported by the DRAM Controller by using the CAS-before-RAS refresh protocol. All four $\overline{\text{DramCAS}}()$ lines are asserted for CASW time followed after 1 clock by asserting the even $\overline{\text{DramRAS}}()$ lines for CASW+0.5 time followed (staggered) by the odd $\overline{\text{DramRAS}}()$ lines asserting for CASW+0.5 clocks time. All four $\overline{\text{DramCAS}}()$ lines, as shown in Figure 10.31, de-assert 1.5 clocks after the odd $\overline{\text{DramRAS}}()$ lines assert. Staggering the RAS lines allows the peak power consumption of turning DRAM chips on to be minimized. The DRAM Controller guarantees that the write enables, $\overline{\text{DramWrEn}}(\text{Odd/Even})$ are de-asserted during refreshes to avoid entry into an internal test mode of higher density (4-16Mbit) DRAM chips.

Refresh cycles can occur in parallel with non-DRAM accesses. If a CPU or DMA transfer requiring DRAM occurs concurrently or after a refresh, the refresh has priority and will complete first.

Because refreshes can happen in parallel with non-DRAM accesses, the R36100 Debug Interface provides a $\overline{\text{DiagNoCS}}()$ signal to decode precisely when a load or store occurs when no chip select (or RAS line) is asserted for a load/store transaction.

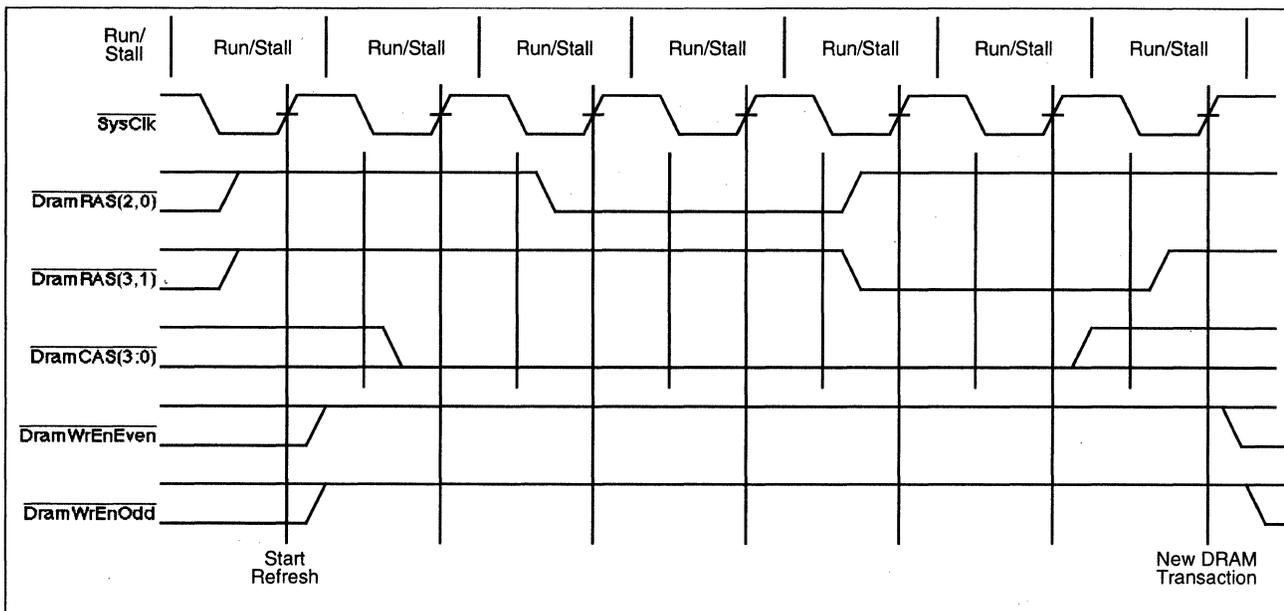


Figure 10.31 DRAM Staggered Refresh

System Examples

The following DRAM systems concentrate on distinguishing the data path connections between the three different DRAM types:

- FCT245 Transceiver Type
- FCT260 Latched Multiplexer (Bus Exchanger) Type
- FCT543 Registered Transceiver Type

The address path of a particular system will depend on the total number of loads the address bus needs to drive. Typically there are 8 DRAM chips per bank, each of them with an address connection. Assuming that a ROM bank is also connected, that is already 12 loads. Assuming that the DRAM and the EPROM are CMOS type input loads (micro-amps), typically the drive current from the R36100 is rarely an issue. However, as the number of loads gets larger, the output propagation delay will also have a capacitive delay factor as well as a noise factor from the trace length. If more than about 8 loads are connected to the SysAddr bus, then allowances in the programmable timing settings of the DRAM Controller should be made for ringing and settling time as well as capacitive load delay derating. If optimal timing is still desired, then address buffers such as the FCT244, FCT344, or FCT827 can be used.

DRAM System using FCT245 Transceivers

DRAM Systems using FCT245 transceivers can be expanded from 1 bank to 4 banks. The first bank, even bank $\overline{\text{DramRAS}}(0)$, uses one set of transceivers and shares the transceiver set with the other optional even bank $\overline{\text{DramRAS}}(2)$. If present, the second bank, odd bank $\overline{\text{DramRAS}}(1)$, uses a separate set of transceivers and shares the transceiver set with the other optional odd bank $\overline{\text{DramRAS}}(3)$. The use of a second set of transceivers allows even and odd banks to be used in the interleaved mode.

In an FCT245 type system, $\overline{\text{DramRdEn}}(\text{Odd,Even})$ are used as the common output enable. Thus $\overline{\text{DramRdEn}}(\text{Odd,Even})$ for the FCT245 type, assert for both reads and writes and could be called, " $\overline{\text{DramEn}}(\text{Odd,Even})$." Because the FCT245 does not contain a latch, address pipelining optimization cannot occur. The use of $\overline{\text{SysRd}}$ or perhaps $\overline{\text{SysWr}}$ (depending on whether the data path from the CPU is A to B or perhaps B to A) avoids leading edge bus contention from direction to output enable skew. Note that the use of $\overline{\text{SysRd}}$ or $\overline{\text{SysWr}}$ on DRAM accesses may in the future, limit the use of the future use of DMA fly by accesses. (The R36100 does not presently support DMA fly by accesses. The other DRAM types described below do not use $\overline{\text{SysRd}}$ or $\overline{\text{SysWr}}$).

Note that the R36100 DRAM Controller depends on $\overline{\text{CAS}}$ without $\overline{\text{RAS}}$ having no effect (always true of standard DRAMs, since some chips do not have dedicated output enable pins) in order to share transceivers between DRAM chip banks.

With present day pricing, the FCT245 type system is the least expensive interleaved option, however, it is not as fast on burst reads as the other two types.

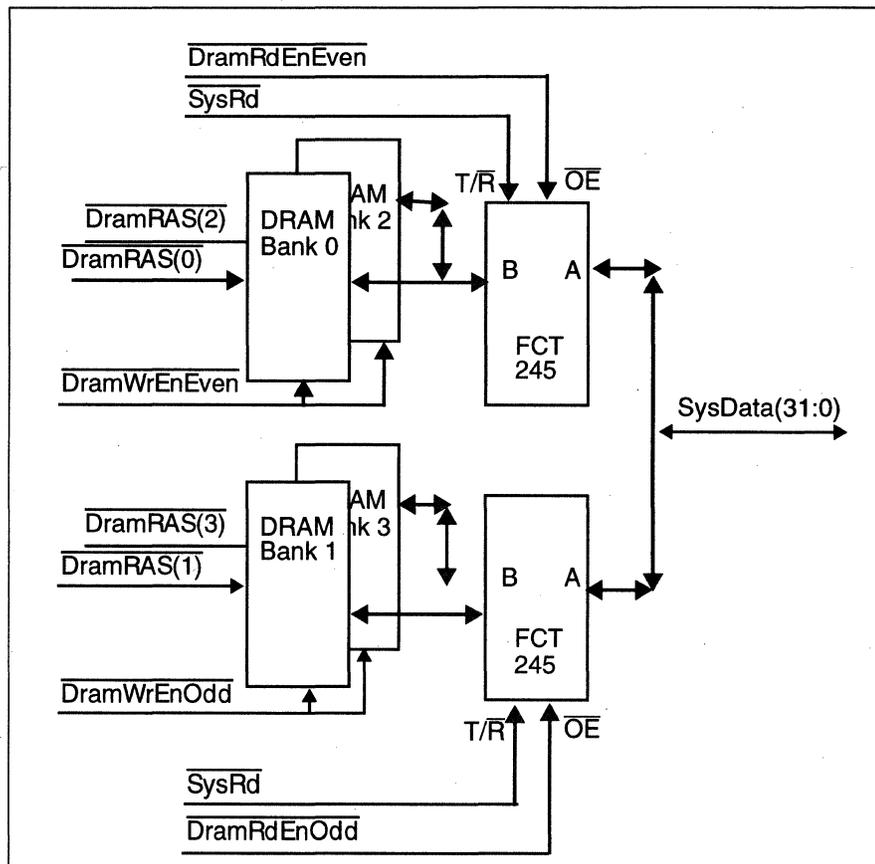


Figure 10.32 Interleaved FCT245 Interface

Low Cost DRAM System using FCT245 Transceivers

In very low cost systems that do not need the extra throughput of interleaving, a single set of transceivers can be used for all 4 banks. However, this requires that the banks not be put into their software programmed interleaved mode and that the read enables, $\overline{\text{DramRdEn}}(\text{Odd,Even})$ be externally OR'ed.

Very Low Cost DRAM System without Transceivers

In simpler systems, it is also possible to remove the transceivers completely, such that the DRAM bank is attached directly to the SysData bus. The Bus Turn Around setting can be adjusted to prevent bus contention between DRAM chips and the CPU on a DRAM read followed by a CPU write. For more information, refer to "Dram Read Cycle Bus Turn-Around ('DramRdBTA') Field Encoding" (Table 10.23 on page 20).

DRAM System using FCT260 Multiplexers

DRAM Systems using a set of FCT260 latched multiplexers can be expanded from 1 bank to 4 banks. The even banks share one data path while the odd banks share the other data path.

In an FCT260 system, $\overline{\text{DramRdEnEven}}$ is used as the common read data path enable to the CPU. $\overline{\text{DramRdEnEven}}$ for the FCT260 type, asserts for both even and odd reads and could be called, " $\overline{\text{DramRdEn}}$." $\overline{\text{DramRdEnOdd}}$ is used to latch the odd read data temporarily so that address pipelining can occur. $\overline{\text{DramRdEnOdd}}$ is also used for the FCT260 path select. The DRAM write enables, $\overline{\text{DramWrEn}}(\text{Odd/Even})$ are hooked up in a straightforward manner, to the odd and even write data path enables of the FCT260, respectively.

The FCT260 system (see Figure 10.33 for system diagram) is one of the least expensive interleaved options, since just 3 chips are required instead of 4 chips. In addition, burst reads are fully optimized with address pipelining, and thus save an additional clock on each burst read relative to a FCT245 system. Thus for many multi-bank systems, the FCT260 system is the best cost/performance alternative.

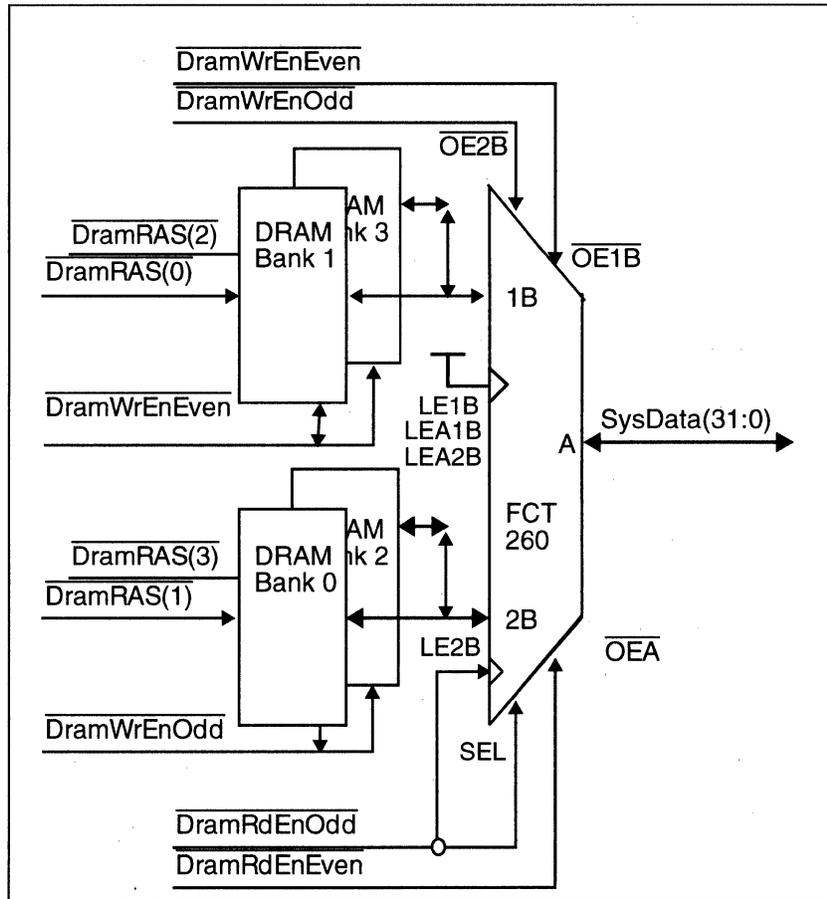


Figure 10.33 Interleaved FCT260 Interface

DRAM System using FCT543 Registered Transceivers

DRAM Systems using a set of FCT543 registered transceivers can be expanded from 1 bank to 4 banks. The even banks share one set of transceivers while the odd banks share another set of transceivers.

The FCT543 system may be more expensive than other options and has the same performance as the FCT260 option. Because the connections are more straightforward, and therefore easier to understand, the FCT543 option is mentioned here as an example system.

In the FCT543 system (see Figure 10.34 for system diagram), the read enables, $\overline{\text{DramRdEn}}(\text{Odd/Even})$, are hooked up to the odd and even transceiver banks' read data path enables respectively. Likewise, the write enables, $\overline{\text{DramWrEn}}(\text{Odd/Even})$, are hooked up to the odd and even transceiver banks' write data path enables respectively.

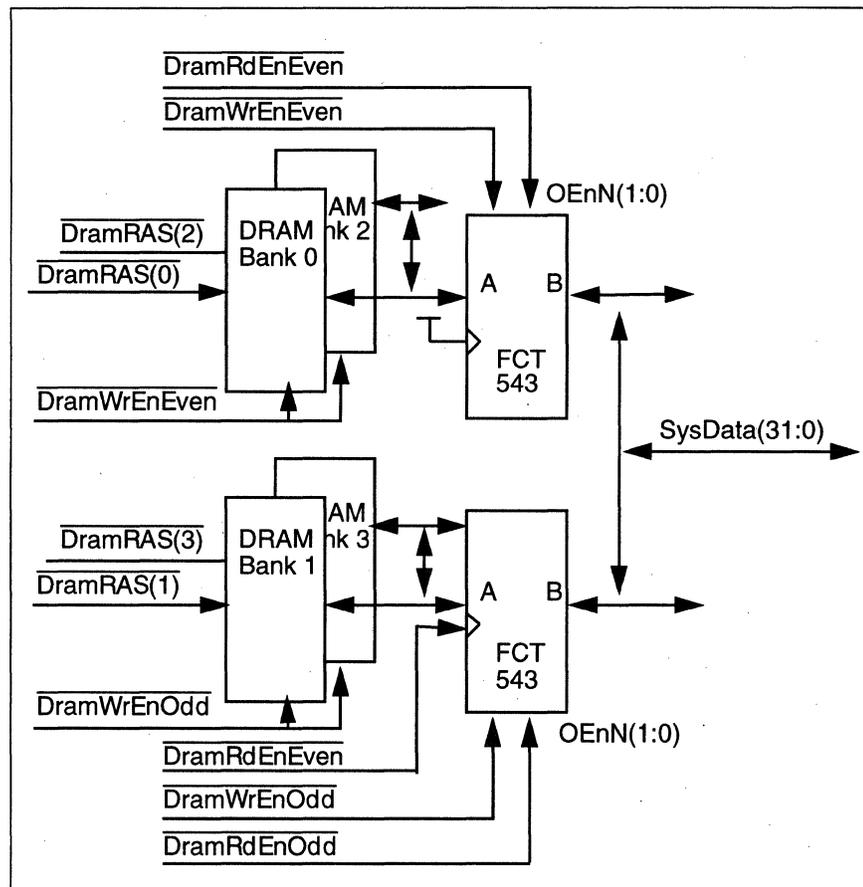


Figure 10.34 Interleaved FCT543 Interface

Block Diagram

The functional block diagram for the address and data path of the DMA Controller is shown in Figure 11.1. The DMA Controller, as one of the bus controllers, coordinates and shares the Bus Interface Unit (BIU) resources with the CPU. An arbitration unit (not shown) coordinates when the DMA Controller Channels can make use of the BIU.

When an internal DMA Controller channel gets the BIU, it first uses the DMA Addr Generator to put a source address out to the BIU. The various BIU control signals, including Endianness and AccTyp (Byte Enables and Burst Length) are also coordinated by the DMA Controller. The BIU then generates a read to the System Interface.

The System Interface executes the read from the source (for instance to DRAM or to one of the on-chip peripherals). If the source read data is of a different port width than 32-bits, then the BIU takes care of the byte gathering, while the Read Buffer takes care of burst gathering. The source data is FIFO'ed into the Read Buffer just like a CPU read.

At this point, the DMA Controller takes the read data from the Read Buffer and generates the target address. The read data is then sent to the BIU in 32-bit quantities with the proper Endianness and AccTyp (Byte Enables and Burst Length) and target address until the Read Buffer is emptied. On each BIU write, the System Interface is invoked and the write is completed out to the target (for instance to one of the on-chip peripherals or to DRAM).

At the same time the DMA Controller is writing to the BIU, it can also optionally simultaneously update the data cache.

Thus the DMA Controller takes control of the Bus Interface Unit and automatically coordinates memory to memory transfers.

Overview

The Direct Memory Access (DMA) controller has two basic functions:

- Internally generated DMA capable of driving slave DMA devices
- External master DMA through internal controllers

Internal DMA Channels

There are four independent DMA Channels in the R36100. Each channel has identical functionality except for their priority encoding. Each channel of the internal DMA Controller is initialized with a set of chaining registers to determine:

- the base start address of the DMA source
- the base start address of the DMA target
- the number of datum to be transferred
- the selection of various protocol styles

Thus the programmable link chaining registers act as a set of instructions to the DMA channel which the channel needs to execute and complete. After completing the link chaining register instructions, the DMA channel may be instructed to stop and interrupt the CPU core, or it may be instructed to load a new set of link chaining register instructions.

At the beginning of a DMA transaction, the channel will first arbitrate for the system bus. By default, all DMA devices and the CPU rotate their priority, as shown in Figure 11.2. When rotating priority is used, the channel with the highest priority will keep its priority privilege until it requests the bus and receives a grant; it will then relinquish the privilege to the next channel in rotation. The exception is for DRAM Refresh which maintains a higher priority than any other rotating channel.

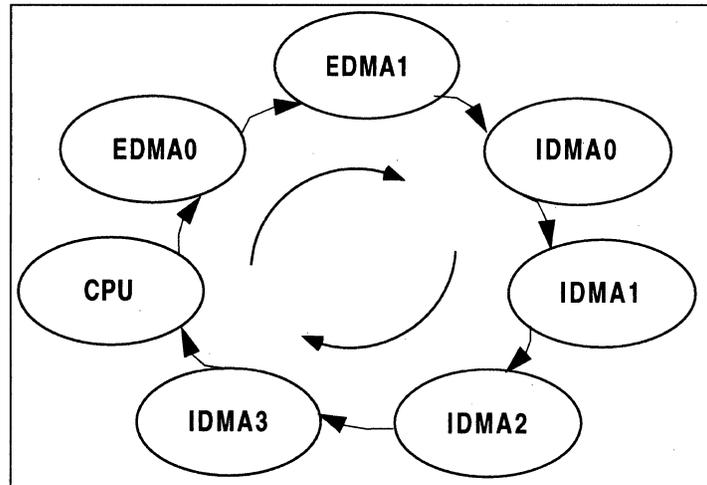


Figure 11.2 Rotating Priority Scheme

Alternatively, the DMA channels can have fixed priority, as shown in Table 11.1.

Fixed Priority	Agent
Highest	DRAM Refresh
•	EDMA0
•	EDMA1
•	IDMA0
•	IDMA1
•	IDMA2
•	IDMA3
Lowest	CPU

Table 11.1 Fixed Priority Encoding

Once arbitration is granted, the DMA channel will generate a read cycle with the source base address. The control register determines whether it is a burst read or not. Typically, the source address will be through an internal memory controller on the R36100 (e.g. the DRAM Controller). Thus the internal memory controller will take the address and generate data, acknowledges, etc., back to the DMA controller channel. The DMA controller uses the BIU 4-word deep buffer FIFO to absorb the potential burst read datum.

After the read is completed, the DMA channel initiates a write to the target address, by emptying out the read buffer FIFO. As with the read, the write is typically through an internal memory controller on the R36100 (e.g. the I/O Controller). The internal memory controller will take the address and data from the FIFO and generate a write transaction.

At the end of the transaction, the DMA channel's count register is decremented by 1. If the count register has not reached 0, the source and target addresses are incremented to their next value (which could be by +0, +1, +2, +4, +8, or +16 in decimal depending on whether incrementing is enabled and whether a mini-burst or burst occurred).

If the count register has reached 0, then the DMA channel is finished with its current link chaining register assignment. If the control register so instructs, the channel may set an interrupt and/or stop, and/or it may reload a new link/set of chaining registers. If a new link/set is loaded, then the DMA channel will repeat the basic DMA channel transaction by copying the new link instructions into the current instructions and executing them.

Internal DMA Algorithm

Figure 11.3 shows the internal DMA algorithm.

```

while (stop_field == false) /* note that if at anytime stop_field == true then break */
{
if (count_field != 0)
/* Start up */
if (wait_for_interrupt_field == true)
{ while (DMAInterruptN == false)
{ /* wait for DMAInterruptN */; }}

assert BusReqN = active_low;
while (BusGntN == non_active_low)
{ /* wait for BusGntN */; }

/* Do the source read */
BusInterfaceUnit(source_addr_field, BurstN,
BEnN(3:0), BigEndianFlag, RdN);
while (Bus_Interface_Unit(FIFO_Data_WrN)) {
FIFO_Data[] = Data;
}
/* Do the target write */
BusInterfaceUnit(target_addr_field, Burst,
BEnN(3:0), BigEndianFlag, WrN);
while (Bus_Interface_Unit(FIFO_Data_RdN))
{ Data = *FIFO_Data++; }

/* Finish up */
if (keep_bus_field == false)
assert BusReqN = non_active_low;
count_field = count_field - 1;
if (source_inc_field == true)
{source_addr_field = source_addr_field + burst_length_field;}
if (target_inc_field == true)
{ target_addr_field = target_addr_field + burst_length_field; }
if (wait_for_interrupt_field == false)
{ break }
} /* if count != 0 */
else { /* count == 0 */
assert ExcInt(DMA_Done_Int()); /* pulse DMADoneN for 1 clock */
if (break_field == true) {
assert Stop_Field == true;
break;
}
else { /* break == false */
case link_field {
0: { DMA_Registers = LinkA_Registers;
break /* from case */; }
1: { DMA_Registers = LinkB_Registers;
break /* from case */; }
2: { DMA_Registers = LinkC_Registers;
break /* from case */; }
3: { DMA_Registers = LinkD_Registers;
break /* from case */; }
}
}
} /* end count == 0 */
} /* end while */

```

Figure 11.3 Internal DMA Algorithm.

External DMA Channels

External DMA channels are conceptually much simpler than their internal counterparts. Much of the control logic is implemented by the external DMA controller agent; so essentially, all the R36100 is required to do is get off the system bus and react to reads and writes to internal controllers.

Thus the R36100 first gives the bus to the external DMA Agent which issues either a read or write command to the R36100. The external DMA Agent then gets off the address and control bus. The R36100 then executes the command on one of the memory controllers and does fly-by data accesses where the external DMA Agent either reads or writes the data at the same time the memory controller writes or reads the data.

The external DMA Controller uses the customary bus request/grant, $\overline{\text{DMABusReq}}()$ and $\overline{\text{DMABusGnt}}()$ handshake signals to allow an external DMA agent to take control of the bus. When the DMA agent takes the bus it drives address and control information onto the pins of the R36100. $\overline{\text{SysRd}}$ is tri-stated; however, its input is ignored since the information is redundant with $\overline{\text{SysWr}}$ being de-asserted. Single word writes need to set the byte enables via $\overline{\text{MemWrEn}}(3:0)$.

Note that if all $\overline{\text{MemWrEn}}()$ signals are left high (de-asserted) the R36100 will automatically interpret this to be equivalent to the all low (asserted) case. Also note that internal peripherals must use single word reads, not burst reads. $\overline{\text{SysALEn}}$ may be asserted for more than one clock; however, the address is latched in on the first rising clock edge where $\overline{\text{SysALEn}}$ is asserted. This allows $\overline{\text{SysALEn}}$ to follow the PCI FRAME# convention. In PCI mode, $\overline{\text{MemWrEn}}(3:0)$ are sampled on the clock after $\overline{\text{SysALEn}}$ is first asserted. The SysData bus is used to get the external DMA 32-bit physical address.

If the physical address corresponds to an on-chip controller and the transfer is a read, the R36100 will generate a read from the proper device and drive the necessary data lines with the data. If the transaction is a burst read or burst write, to indicate the end of the transaction, the accesses must be full word reads or writes and use a properly programmed $\overline{\text{DMADone}}$ input. $\overline{\text{SysWait}}$ is available during external DMA and can be an inverted input and can be used to slow down reads or writes.

Because the R36100 must reuse the SysAddr and Sys Control lines during the second half of a DMA access, the DMA agent must tri-state its address (and data transceivers) after driving SysAddr, $\overline{\text{SysALEn}}$, $\overline{\text{SysBurstFrame}}$, $\overline{\text{SysWr}}$, and $\overline{\text{MemWrEn}}(3:0)$ into the R36100. The tri-stating must occur by the second clock after de-asserting $\overline{\text{SysALEn}}$. Before taking the bus back, the CPU drives all control lines de-asserted; therefore, the DMA agent doesn't necessarily have to do so. With the exception of write SysData(), the R36100 will take over the bus to do a memory cycle. When the R36100 has completed its internal memory cycle, it will assert $\overline{\text{SysDataRdy}}$ and de-assert SysAddr and all System Control lines.

Burst transfers may be from 1 to 64 words and must be aligned with a 64-word block. Burst transfer mode also requires the use of the $\overline{\text{DMADone}}$ pin. $\overline{\text{DMABusReq}}()$ must be kept asserted at least one clock after $\overline{\text{DMABusGnt}}()$ asserts. $\overline{\text{DMABusReq}}()$ must be deasserted before the last $\overline{\text{SysDataRdy}}$ occurs, unless another external DMA transaction is to occur.

Pin Descriptions

Direct Memory Access (DMA) Controller Signals

DmaBusReq(1:0)

Input

DMA Bus Request: Active low input which signals the R36100 that the external DMA controller would like to gain mastership of the system bus. DmaBusReq can be software programmed to be active high by using the ReqH field in the External DMA Control Register.

DmaBusGnt(1:0)

Output

DMA Bus Grant: Active low output which signals to the external DMA controller that it is now master of the system bus. DMABusGnt can be software programmed to be active high by using the ReqH field in the External DMA Control Register.

DmaDone

Input

DMA Done: Active low signals the R36100 that the current DMA transaction is the last transaction by the current DMA agent.

For internal DMA, if DMADone is asserted the present link will abort. For external DMA, DMADone will indicate that the present datum is the last datum on a burst read or write. DMADone is required for burst transfers. It must assert at least two clocks before the last clock cycle of the last datum. Thus it is recommended that burst transfers use at least 2 clocks for each datum. (If 1 clock is used for each datum, then DMADone must be asserted with the next to last datum instead of with the last datum).

System Control Signals used during DMA Controller Accesses

<u>SysALEn</u>	Output/(Input during External DMA)
<u>SysBurstFrame</u>	Output/(Input during External DMA)
<u>SysRd</u>	Output/(Input during External DMA)
<u>SysWr</u>	Output/(Input during External DMA)
<u>MemWrEn(3:0)</u>	Output/(Input during External DMA)

These signals are used during the initial part of an external DMA access, in order to give the R36100 a read, write, burst read, or burst write command. During this period, they are inputs.

SysRd:

SysRd can optionally be driven by the external DMA agent; however, it is ignored by the R36100, which uses an unasserted SysWr to indicate a read command.

MemWrEn():

External DMA, in addition to the regular cases, can support the '1111' case which is equivalent to the '0000' all asserted case and the '1001' case which is a case the R3000A core never generates.

SysDataRdy**Output**

SysDataRdy asserts low whenever the CPU expects data to be read or written. It is always a CPU output. The datum is read or written by the CPU bus controller simultaneous with being written/read by the DMA agent.

Register Descriptions

The R36100 DMA Controller has two sets of registers, one internal (4 channels) and one external (2 channels).

Internal DMA Controller Register Descriptions

Table 11.2 is an address map of the Internal DMA Controller registers. Big Endian software must offset these addresses by b'10 (0x2).

Physical Addr	Description
0xFFFF_E300	DMA LSB Source Address Register for Channel 0
0xFFFF_E304	DMA MSB Source Address Register for Channel 0
0xFFFF_E308	DMA LSB Target Address Register for Channel 0
0xFFFF_E30C	DMA MSB Target Address Register for Channel 0
0xFFFF_E310	DMA LSB Count Register for Channel 0
0xFFFF_E314	DMA MSB Count Register for Channel 0
0xFFFF_E318	DMA LSB Control Register for Channel 0
0xFFFF_E31C	DMA MSB Control Register for Channel 0
0xFFFF_E320	DMA LSB Source Address Register for Channel 1
0xFFFF_E324	DMA MSB Source Address Register for Channel 1
0xFFFF_E328	DMA LSB Target Address Register for Channel 1
0xFFFF_E32C	DMA MSB Target Address Register for Channel 1
0xFFFF_E330	DMA LSB Count Register for Channel 1
0xFFFF_E334	DMA MSB Count Register for Channel 1
0xFFFF_E338	DMA LSB Control Register for Channel 1
0xFFFF_E33C	DMA MSB Control Register for Channel 1
0xFFFF_E340	DMA LSB Source Address Register for Channel 2
0xFFFF_E344	DMA MSB Source Address Register for Channel 2
0xFFFF_E348	DMA LSB Target Address Register for Channel 2
0xFFFF_E34C	DMA MSB Target Address Register for Channel 2
0xFFFF_E350	DMA LSB Count Register for Channel 2
0xFFFF_E354	DMA MSB Count Register for Channel 2
0xFFFF_E358	DMA LSB Control Register for Channel 2
0xFFFF_E35C	DMA MSB Control Register for Channel 2
0xFFFF_E360	DMA LSB Source Address Register for Channel 3
0xFFFF_E364	DMA MSB Source Address Register for Channel 3
0xFFFF_E368	DMA LSB Target Address Register for Channel 3
0xFFFF_E36C	DMA MSB Target Address Register for Channel 3
0xFFFF_E370	DMA LSB Count Register for Channel 3
0xFFFF_E374	DMA MSB Count Register for Channel 3
0xFFFF_E378	DMA LSB Control Register for Channel 3
0xFFFF_E37C	DMA MSB Control Register for Channel 3
0xFFFF_E380	DMA LSB Source Address Register for Link A
0xFFFF_E384	DMA MSB Source Address Register for Link A
0xFFFF_E388	DMA LSB Target Address Register for Link A
0xFFFF_E38C	DMA MSB Target Address Register for Link A
0xFFFF_E390	DMA LSB Count Register for Link A
0xFFFF_E394	DMA MSB Count Register for Link A
0xFFFF_E398	DMA LSB Control Register for Link A
0xFFFF_E39C	DMA MSB Control Register for Link A
0xFFFF_E3A0	DMA LSB Source Address Register for Link B
0xFFFF_E3A4	DMA MSB Source Address Register for Link B
0xFFFF_E3A8	DMA LSB Target Address Register for Link B
0xFFFF_E3AC	DMA MSB Target Address Register for Link B
0xFFFF_E3B0	DMA LSB Count Register for Link B
0xFFFF_E3B4	DMA MSB Count Register for Link B
0xFFFF_E3B8	DMA LSB Control Register for Link B
0xFFFF_E3BC	DMA MSB Control Register for Link B
0xFFFF_E3C0	DMA LSB Source Address Register for Link C
0xFFFF_E3C4	DMA MSB Source Address Register for Link C
0xFFFF_E3C8	DMA LSB Target Address Register for Link C
0xFFFF_E3CC	DMA MSB Target Address Register for Link C
0xFFFF_E3D0	DMA LSB Count Register for Link C
0xFFFF_E3D4	DMA MSB Count Register for Link C
0xFFFF_E3D8	DMA LSB Control Register for Link C
0xFFFF_E3DC	DMA MSB Control Register for Link C

Table 11.2 Internal Channel DMA Controller Register Address Map

DMA LSB Source Address Register for Channel 0..3
 ('DmaLSBSourceAddrReg(0..3)')

DMA LSB Source Address Register for Link A..D
 ('DmaLSBSourceAddrReg(A..D)')

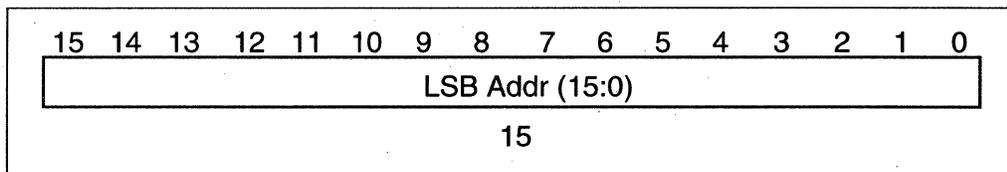


Figure 11.4 Internal DMA LSB Source Address Register ('DmaLSBSourceAddrReg').

DMA MSB Source Address Register for Channel 0..3
 ('DmaMSBSourceAddrReg(0..3)')

DMA MSB Source Address Register for Link A..D
 ('DmaMSBSourceAddrReg(A..D)')

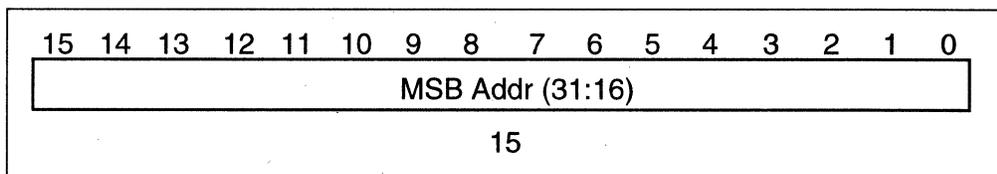


Figure 11.5 Internal DMA MSB Source Address Register ('DmaMSBSourceAddrReg').

The Source Address Register, shown in Figure 11.5, must be programmed with the initial address of the peripheral or memory that data is to be read from. The channel Source Address Register will be incremented by the Source Burst Size amount after each DMA transaction. Normally the Source Address Register is only written; however, it may also be read for diagnostic reasons.

DMA LSB Target Address Register for Channel 0..3
 ('DmaLSBTargetAddrReg(0..3)')

DMA LSB Target Address Register for Link A..D
 ('DmaLSBTargetAddrReg(A..D)')

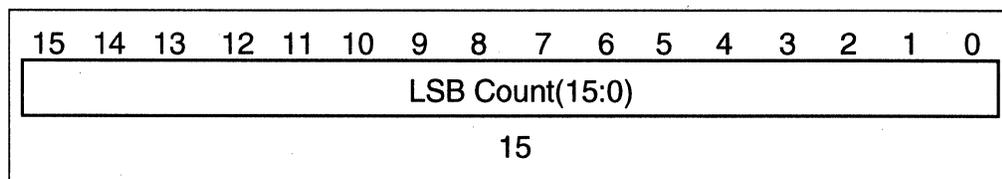


Figure 11.6 Internal DMA LSB Target Address Register ('DmaLSBTargetAddrReg').

DMA MSB Target Address Register for Channel 0..3
 ('DmaMSBTargetAddrReg(0..3)')

DMA MSB Target Address Register for Link A..D
 ('DmaMSBTargetAddrReg(A..D)')

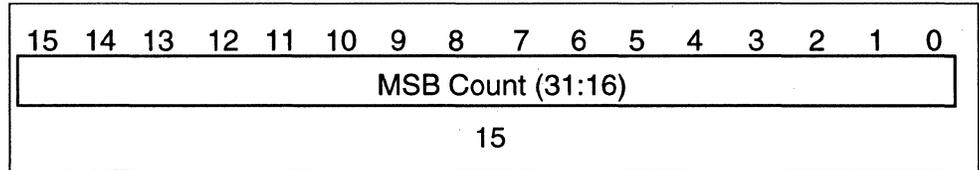


Figure 11.7 Internal DMA MSB Target Address Register ('DmaMSBTargetAddrReg').

The Target Address Register (shown in Figure 11.6 and Figure 11.7) must be programmed with the initial address of the peripheral or memory that data is to be written to. The channel Target Address Register will be incremented by the Target Burst Size amount after each DMA transaction. Normally the Target Address Register is only written, however, it may also be read for diagnostic reasons.

DMA LSB Count Register for Channel 0..3
 ('DmaLSBCountReg(0..3)')

DMA LSB Count Register for Link A..D
 ('DmaLSBCountReg(A..D)')

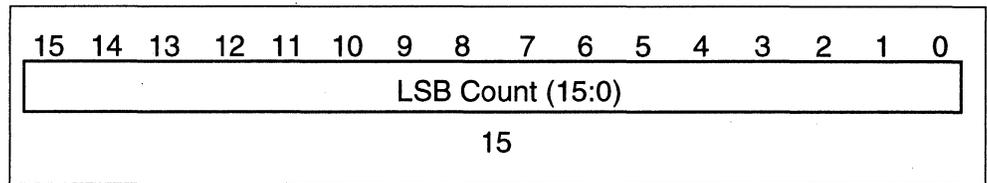


Figure 11.8 DMA LSB Count Register ('DmaLSBCountReg').

DMA MSB Count Register for Channel 0..3
 ('DmaMSBCountReg(0..3)')

DMA MSB Count Register for Link A..D
 ('DmaMSBCountReg(A..D)')

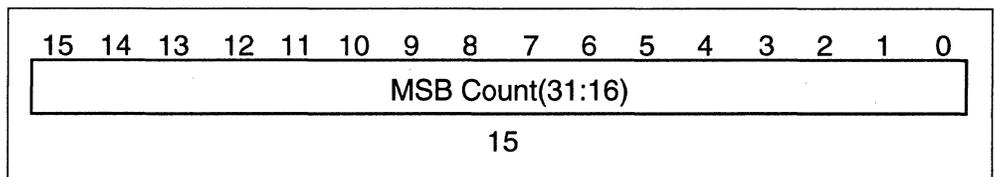


Figure 11.9 Internal DMA MSB Count Register ('DmaMSBCountReg').

Note: The count is the number of read/write transactions to be done. A burst or mini-burst read/write only counts once. The channel count register is decremented by 1 after each DMA transaction. If the count is 0, then the DMA Channel will immediately proceed to the next link. The LSB and MSB registers are writable. They are also readable for diagnostic purposes.

The InternalDMA LSB Control register is shown in Figure 11.10, with bit assignments listed in Table 11.3.

DMA LSB Control Register for Channel 0..3
 ('DmaLSBControlReg(0..3)')

DMA LSB Control Register for Link A..D
 ('DmaLSBControlReg(A..D)')

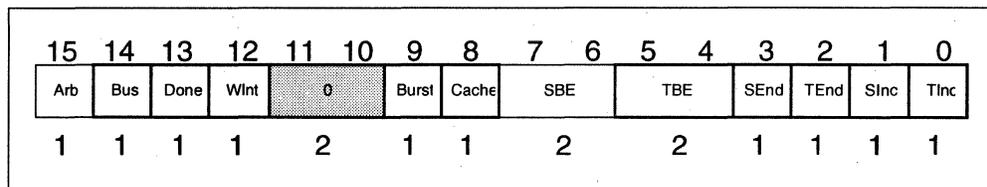


Figure 11.10 Internal DMA LSB Control Register ('DmaLSBControlReg').

Bit	Assignment
15	Arbitration Type
14	Keep Bus
13	Allow DMA Done
12	Wait for Interrupt
11:10	'0'
9	Burst Type
8	Cacheable Type, Access Type(2)
7:6	Source Byte Enable Type, Access Type(1:0)
5:4	Target Byte Enable Type, Access Type(1:0)
3	Source Endianness
2	Target Endianness
1	Increment Source
0	Increment Target

Table 11.3 Internal DMA LSB Control Register ('DmaLSBControlReg') Bit Assignments.

Arbitration Type ('Arb') Field:

The Arbitration Type is only applicable to Channel 0. However, all channels are required to be programmed to the same value. The Arbitration Type field values are listed in Table 11.4.

Value	Action
'1'	Fixed Priority Arbitration (default)
'0'	Rotating Priority Arbitration.

Table 11.4 Arbitration Type ('Arb') Field Encoding

Keep Bus ('Bus') Field:

Value	Action
'1'	Keep Bus until done with current link.
'0'	Let go of Bus at the end of each read/write.

Table 11.5 Keep Bus ('Bus') Field Encoding

Keep Bus field values are listed in Table 11.5.

Allow DMADone ('Done') Field:

Note: If 'Done' is programmed as a '1' (see Table 11.6 for field encodings), then it is required to program the DMADone pin as an input. Programming the value to '1' forces DMADone to behave as a "stop-link" indicator; if DMADone gets asserted during a channel's DMA BusGnt asserted period, then it will abort the DMA channel and its link after this current transaction completes. Field values and descriptions are listed in Table 11.6.

Value	Action
'1'	"Stop-link"
'0'	Normal.

Table 11.6 Allow DMADone ('Done') Field Encoding.**Wait for Interrupt ('WInt') Field:**

The DMA Controller cannot autonomously acknowledge the source of the interrupt. Thus, it is expected that interrupts either pulse low for 1 clock or self reset when the pertinent data port is read or written.

In the "Wait" mode, the DMA Controller does not sample for the interrupt again until the clock after the internal bus grant is released. If programmed for the continuous mode, the DMA Controller does not sample until the clock after the internal bus grant would have been released. This gives adequate time for the external I/O device being read or written to reset its internal interrupt generator. Refer to Table 11.7 for field values.

Value	Action
'1'	Wait after each transfer until the next interrupt
'0'	Continuous Style.

Table 11.7 Wait for Interrupt ('Int') Field Encoding**Burst Type ('Burst') Field:**

Value	Action
'1'	Transaction is a burst of more than 4 bytes.
'0'	Transaction is of 4 bytes or less.

Table 11.8 Burst Type ('Burst') Field Encoding.

Refer to Table 11.8 for Burst Type field values and descriptions.

Cacheable ('Cache') Field:

If the Cacheable Field is set, then the internal CPU is guaranteed to halt. Otherwise, the CPU may be able to execute instructions out of cache up to the time it writes four words to the write buffer. The CPU write buffer is flushed before DMA can start. Refer to Table 11.9 for Cacheable field values and descriptions.

Value	Action
'1'	Transaction is write cached.
'0'	Transaction is write uncached.

Table 11.9 Cacheable ('Cache') Field Encoding

Source Byte Enable Type ('SBE') Field:

The selected type is combined with Addr(1:0) and Endianness to form the byte-enables. Refer to Table 11.10 for Source Byte Enable field values and descriptions.

Value	Selected Type
'11'	Word.
'10'	Reserved.
'01'	Halfword.
'00'	Byte.

Table 11.10 Source Byte Enable Type ('SBE') Field Encoding.

Target Byte Enable Type ('TBE') Field:

The selected type is combined with Addr(1:0) and Endianness (see Table 11.12 for Source BigEndianness type field values and descriptions or Table 11.13 for Target Endianness values and descriptions) to form the byte-enables. Refer to Table 11.11 for Target Byte field values and descriptions.

Value	Selected Type
'11'	Word.
'10'	Reserved.
'01'	Halfword.
'00'	Byte.

Table 11.11 Target Byte Enable Type ('TBE') Field Encoding

Source Endianness Type ('SEndian') Field:

Value	Endianness
'1'	Big Endian.
'0'	Little Endian.

Table 11.12 Source Big Endianness Type ('SEndian') Field Encoding

Target Endianness Type ('TEndian') Field:

Value	Endianness
'1'	Big Endian.
'0'	Little Endian.

Table 11.13 Target Big Endianness Type ('TEndian') Field Encoding

Increment Source Address ('SInc') Field:

Value	Action
'1'	Increment source address
'0'	Constant source address

Table 11.14 Increment Source Address ('HInc') Field Encoding

Increment Target Address ('TInc') Field:

Value	Action
'1'	Increment target address
'0'	Constant target address

Table 11.15 Increment Target Address ('TInc') Field Encoding

Programming information for the Increment Source and Increment Target fields are listed in Table 11.14 and Table 11.15.

DMA MSB Control Register for Channel 0..3 ('DmaMSBControlReg(0..3)')

DMA MSB Control Register for Link A..D ('DmaMSBControlReg(A..D)')

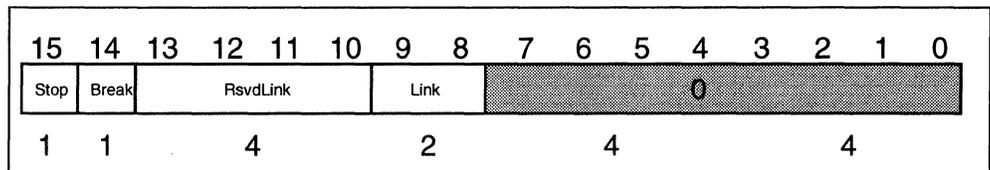


Figure 11.11 Internal DMA MSB Control Register ('DmaMSBControlReg').

The Internal DMA MSB Control register fields are shown in Figure 11.11, with bit assignments listed in Table 11.16. Field values and descriptions are listed in Table 11.17, Table 11.18, Table 11.19, and Table 11.20.

Bit #	Field Name
15	Stop
14	Break
13:10	Reserved
9:8	Link
7:0	Reserved

Table 11.16 Internal DMA MSB Control Register ('DmaMSBControlReg') Bit Assignments

Stop ('Stop') Field:

Value	Action
'1'	Stop Immediately and abort any DMA link in progress (default). Any DMA bus transaction may complete.
'0'	Enable DMA.

Table 11.17 Stop ('Stop') Field Encodin

Break ('Break') Field:

Value	Action
'1'	Break at the end of this DMA chain (default). All Reserved Link and all Link Bits must also be set to 1.
'0'	Execute next link at the end of this DMA chain.

Table 11.18 Break ('Break') Field Encoding

ReservedLink Field ('RsvdLink'):

Value	Action
'1111'	Must be written with the same value as the Break field. Undefined during reads.
'0000'	Must be written with the same value as the Break field. Undefined during reads.

Table 11.19 Reserved Link ('RsvdLink') Field Encoding

Link ('Link') Field:

Value	Action
'3'	Load LinkD at the end of this DMA chain and execute (default).
'2'	Load LinkC at the end of this DMA chain and execute.
'1'	Load LinkB at the end of this DMA chain and execute.
'0'	Load LinkA at the end of this DMA chain and execute.

Table 11.20 Link ('Link') Field Encoding

External DMA Controller Register Descriptions

Table 11.21 is an address map of the External DMA Controller registers. Note that Big Endian software must offset these addresses by b'10 (0x2). The External DMA Control register 0..1 fields are shown in Figure 11.12, bit assignments are listed in Table 11.22. Field values and descriptions are given in Table 11.23, Table 11.24, Table 11.25, and Table 11.26.

Timing Diagrams for External DMA Single Datum Read using the Memory Controller (Figure 11.13 on page 19), External DMA Single Datum Write using Memory Controller (Figure 11.14 on page 20), External DMA Two-Datum Burst Read using the Memory Controller (Figure 11.15 on page 21), and External DMA Two-Datum Burst Write using the Memory Controller (Figure 11.16 on page 21) are also included in this section.

Phys. Addr	Description
0xFFFF_E400	ExtDMA Control Register 0
0xFFFF_E410	ExtDMA Control Register 1

Table 11.21 External DMA Controller Register Address Assignments

External DMA Control Register 0..1 (‘ExtDmaControlReg(0..1)’)

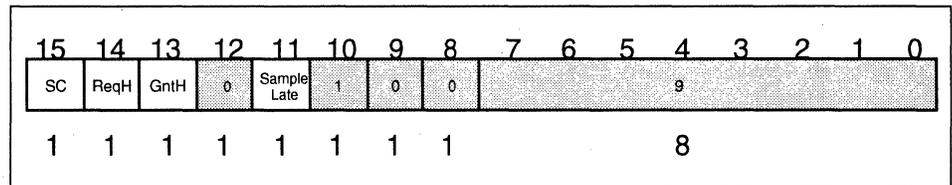


Figure 11.12 External DMA Control Register (‘ExtDmaControlReg’)

Bit	Assignment
15	Stop Channel
14	DMABusReq() active High
13	DMABusGnt() active High
12	Reserved
11	Sample MemWrEn(3:0) and SysBurstFrame one clock later
10	Reserved to ‘1’
9..0	Reserved to ‘0’

Table 11.22 External DMA Control Register (‘ExtDmaControlReg’) Bit Assignments

Stop Channel ('EC') Field:

Value	Action
'1'	Stop/Disable External DMA Channel (default).
'0'	Enable External DMA Channel.

Table 11.23 Enable Channel ('EC') Field Encoding**Bus Request Protocol High ('ReqH') Field:**

Value	Action
'1'	Active High Protocol; $\overline{\text{DMABusReq}}$ is active high for this channel (SCSI controller convention).
'0'	Active Low Protocol; $\overline{\text{DMABusReq}}$ is active low for this channel (default).

Table 11.24 Bus Request Protocol High ('ReqH') Field Encoding**Bus Grant Protocol High ('GntH') Field:**

Value	Action
'1'	Active High Protocol (I-Style convention).
'0'	Active Low Protocol (default).

Table 11.25 Bus Grant Protocol High ('GntH') Field Encoding**Sample $\overline{\text{MemWrEn}}$ and $\overline{\text{SysBurstFrame}}$ 1 Clock Later ('SampleLate') Field:**

This field selects whether to sample $\overline{\text{MemWrEn}}(3:0)$ and $\overline{\text{SysBurstFrame}}$ on writes with $\overline{\text{SysALEn}}$ or one clock later as would be done for PCI accesses.

Value	Action
'1'	Sample $\overline{\text{MemWrEn}}$ and $\overline{\text{SysBurstFrame}}$ one clock after $\overline{\text{SysALEn}}$ first asserts.
'0'	Sample $\overline{\text{MemWrEn}}$ and $\overline{\text{SysBurstFrame}}$ with $\overline{\text{SysALEn}}$.

Table 11.26 Sample $\overline{\text{MemWrEn}}$ and $\overline{\text{SysBurstFrame}}$ 1 clock later ('SampleLate') Field Encoding

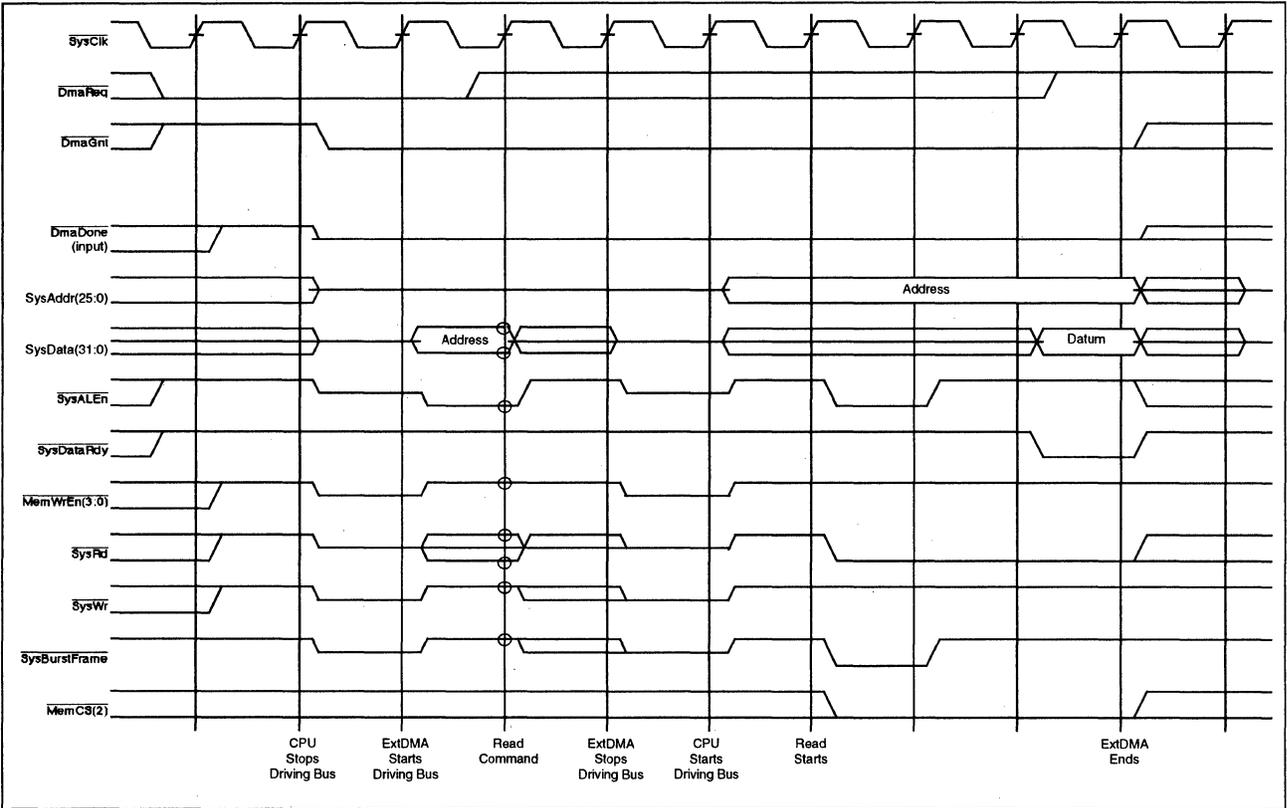


Figure 11.13 External DMA Single Data Read using the Memory Controller (Data Transfer from Memory to Device)

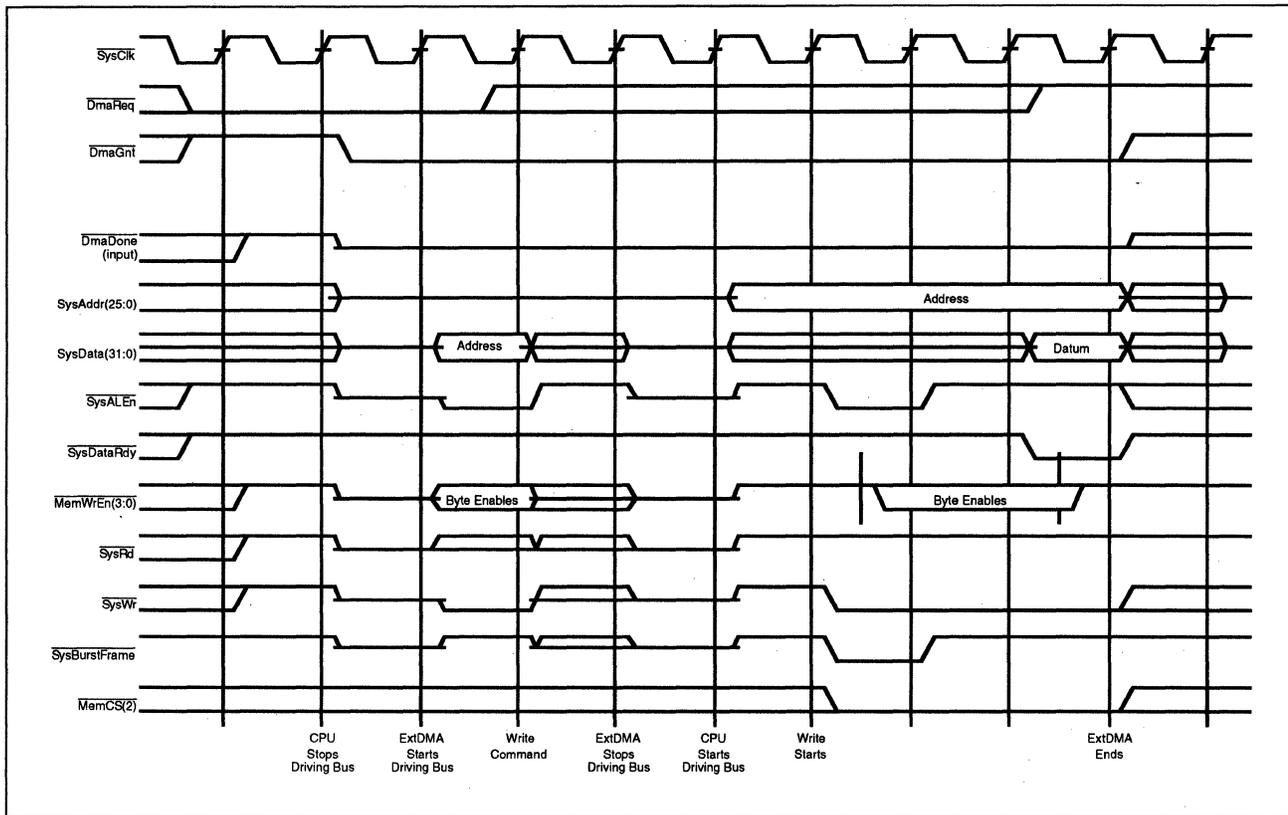


Figure 11.14 External DMA Single Data Write using the Memory Controller (Data Transfer from Device to Memory)

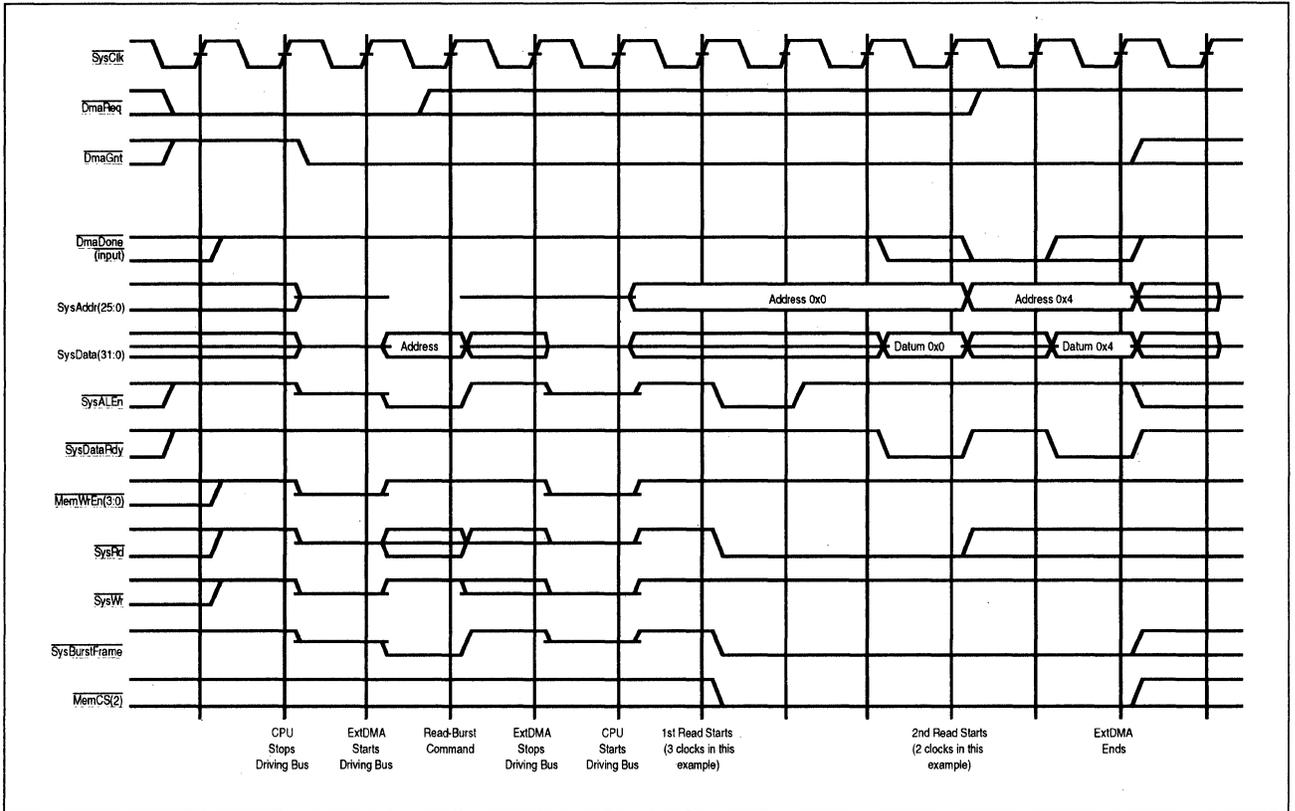


Figure 11.15 External DMA Two-Data Burst Read using the Memory Controller (Data Transfer from Memory to Device)

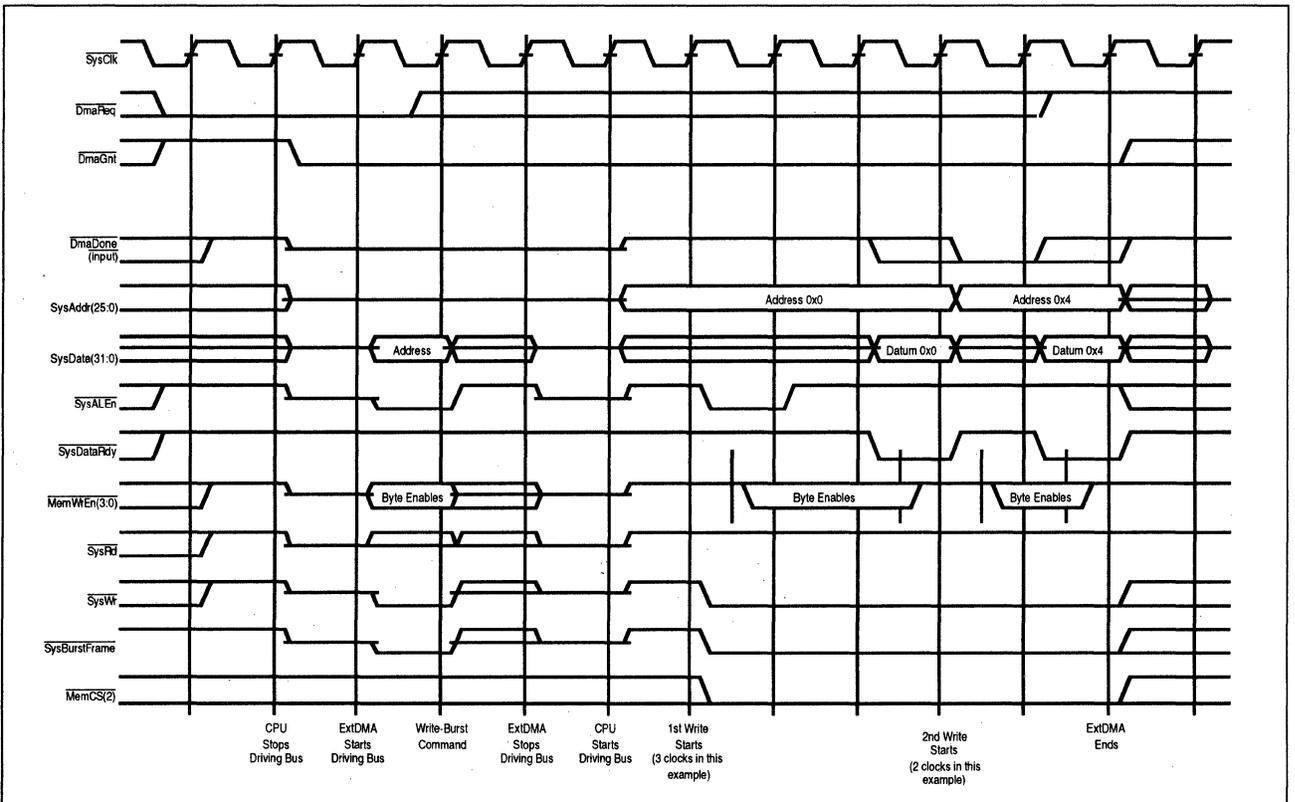


Figure 11.16 External DMA Two-Data Burst Write using the Memory Controller (Data Transfer from Device to Memory)

System Examples

Memory-to-memory Copying

One common system operation may be a memory-to-memory transfer. For example, the system may desire the transfer is from one block of DRAM starting at virtual address 0x80030000 to another block of DRAM starting at virtual address 0x80038000. The block is 1520 bytes long. Only one DMA Channel is needed.

One of the DMA Channels can be programmed so that the source address register is 0x00030000 (the equivalent physical address for virtual address 0x80030000) and target address register 0x00038000. The count register is $1520 / 16 = 95$ since the maximum DMA burst transaction can be up to 16 bytes (the size of the BIU read buffer). In the LSB Control Register, the DMA Channel is set up to Burst, Increment Source address, and Increment Target address. The MSB Control Register is set up to Break after all bytes have been transferred, to have a Burst length for 16 bytes, and to Enable the DMA Channel to begin. The CPU can then continue on with another process while the memory to memory transfer takes place.

Transfers between I/O and Memory

One example where this may be desirable is in copying data between the UART channels and memory. For example, the serial port may be used in full-duplex mode. Thus the serial port is both receiving and transmitting data at the same time. In this scenario, two DMA Channels are used, one for receiving, and one for transmitting.

Using the Interrupt Controller, certain interrupts can be steered to certain DMA Channels. For instance, the Serial Receive Interrupt can be steered to DMA Channel 1 and the Transmit Interrupt can be steered to DMA Channel 0.

In the receiving case, the DMA Channel is programmed to transfer data from the serial port to a DRAM buffer block. When the buffer becomes full, the CPU can act on a DMA Channel 1 Done Interrupt and/or set up another DRAM buffer block via the link registers. The DMA Channel is programmed to not Increase the Source Address for the serial port and to Increase the Target Address after each byte transfer. The DMA Channel is set to transfer 1 byte per count (for the burst length) and to Wait for Interrupt between bytes. After each byte is transferred, the DMA Channel will wait for another receive interrupt before transferring another byte from the serial port to DRAM.

In the transmit case, the DMA Channel is programmed to transfer the message length number of bytes from DRAM which Increases the Source Address after each byte to the Serial Port which does not Increase the Target Address after each byte. The DMA Channel is set to transfer 1 byte per count (as the burst length) and to Wait for Interrupt between bytes. (Alternatively, since the serial port's transmit FIFO is two bytes, DMA could transfer two bytes at a time). After each byte is transferred, the DMA Channel will wait for another transmit FIFO empty interrupt before transferring another byte from DRAM to the serial port.

By using the two DMA Channels, the CPU is freed from either having to constantly poll for serial port status or from constantly handling interrupts and interrupt service routines.

Distinguishing Between CPU and Internal DMA Accesses

If the external system needs to distinguish between a bus transaction generated by the CPU core versus one of the internal DMA channels, there are a couple of options:

- The system software can code an internal MSB address for the internal DMA channel, and have that address bit be ignored by the address map. The assertion of this address would then signal a DMA transfer.
- Use the DIAGInternalDMA pin.



Introduction

The IDT79R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. The on-chip peripherals include **Parallel Input/Output (PIO) Pins** (see Figure 12.1 for block diagram) as described in this chapter.

This chapter provides an overview of the PIO programming interface, a complete description of the signal pins, and discusses how PIOs relate to typical internal and external systems.

Features

- 39 general purpose PIO parallel input/output pins
- PIO pins multiplexed with controller functions

Block Diagram

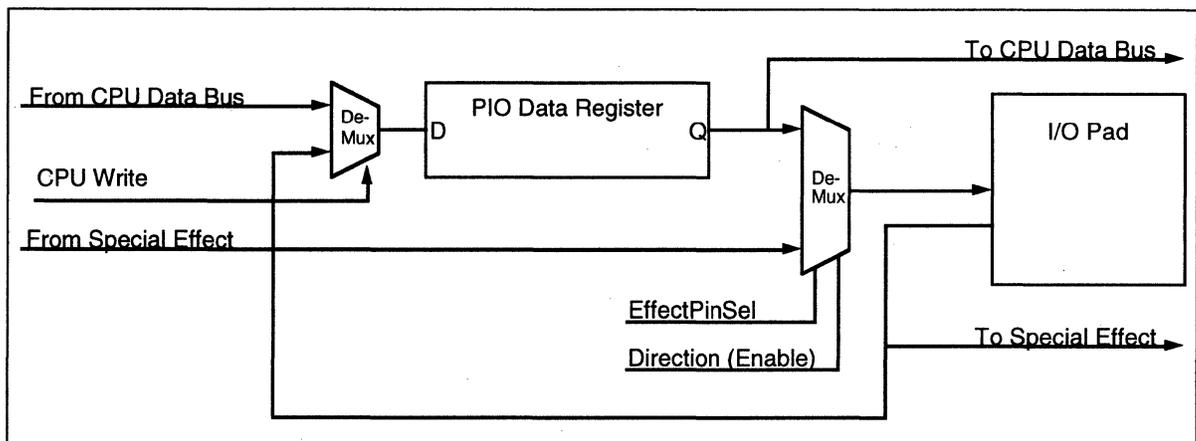


Figure 12.1 PIO Block Diagram.

Overview

The Parallel Input/Output (PIO) pins are multi-purpose pins that can be programmed to act as inputs or outputs. Each PIO pin is also multiplexed with other controller's inputs or outputs. This flexible arrangement allows the system designers to customize the R36100's resources according to their needs. Designs needing a special purpose controller, such as the Laser Video Controller, can allocate the Laser Video pins for that purpose. Other applications, such as Datacom, that do not need the Laser Video Controller can use those pins for general purpose inputs or outputs.

Inputs are not synchronized beyond the requirements of the destination unless otherwise noted. Outputs are non-synchronized (typically they are synchronized by the originating peripheral) and are multiplexed.

Pin Descriptions

Each of the PIO pins is multiplexed with other pin functions. Some PIOs are input only, some are output only, and some can be programmed as either inputs or outputs.

PIO(n) Input/Output

Parallel Input/Output: These bi-directional signals can be used as generic input/output pins. They are set individually through control registers in the PIO interface and can be read by software reads to the appropriate registers.

Table 12.1 shows the relationship between the PIO pins and the other R36100 function pins that are multiplexed onto the same device pin.

PIO Number	Alternate Function	Register Number	Bit Position
PIO(31)	SerialCTS(0)	1	14
PIO(30)	SerialDCD(0)	1	13
PIO(29)	SerialRxData(1)	1	12
PIO(28)	SerialPClkIn	1	11
PIO(27)	SerialCTS(1)	1	10
PIO(26)	SerialDCD(1)	1	9
PIO(25)	LaserVideoClkIn	1	8
PIO(24)	LaserLineSync	1	7
PIO(23)	LaserPageSync	1	6
PIO(22)	CentStrobe	1	5
PIO(21)	CentAutoFeed	1	4
PIO(20)	CentInit	1	3
PIO(19)	CentSelect	1	2
PIO(18)	DMABusReq(1)	1	1
PIO(17)	ExcInt(4)	1	0
PIO(16)	ExcInt(3)	2	14
PIO(15)	BrCond(3)	2	13
PIO(14)	BrCond(2)	2	12
PIO(41)	SerialTxData(0)	2	11
PIO(40)	SerialRxData(0)	2	10
PIO(39)	SerialPClkIn(0)	2	9
PIO(38)	SerialSClk(0)	2	8
PIO(37)	SerialSync(0)	2	7
PIO(36)	SerialSClk(1)	2	6
PIO(35)	SerialSync(1)	2	5
PIO(34)	TimerTc/Gate(2)	2	4
PIO(33)	TimerTc/Gate(1)	2	3
PIO(32)	TimerTc/Gate(0)	2	2
PIO(13)	SerialRTS(0)	0	13
PIO(12)	SerialDTR(0)	0	12
PIO(11)	SerialTxData(1)	0	11
PIO(10)	SerialRTS(1)	0	10
PIO(9)	SerialDTR(1)	0	9
PIO(8)	LaserVideoData	0	8
PIO(7)	CentAck	0	7
PIO(6)	CentBusy	0	6
PIO(5)	CentPError	0	5
PIO(4)	CentSelect	0	4
PIO(3)	CentFault	0	3
PIO(2)	CentHostOEn	0	2
PIO(1)	CentHostStrobe	0	1
PIO(0)	DMABusGnt(1)	0	0

Note: Input-only and output-only refer to the functionality of the pin, when the alternate function is used. The Register-Number and Bit-Position fields describe which of the PIO Data, Direction, and Effect-Select register/bit combinations control that PIO signal.

Table 12.1 Alternate R36100 functions mapped to PIO pins

Register Definitions

Note that Big Endian software must offset these addresses by b'10 (0x2).

Table 12.2 provides an address map and descriptions of the PIO Registers. And Figure 12.2 shows the PIO Data Registers. Additional programming information is located in Table 12.3.

Phys. Address	Description
0xFFFF_EA00	PIO Data Register 0
0xFFFF_EA04	PIO Direction Control Register 0
0xFFFF_EA08	PIO Effect Select Control Register 0
0xFFFF_EA10	PIO Data Register 1
0xFFFF_EA14	PIO Direction Control Register 1
0xFFFF_EA18	PIO Effect Select Control Register 1
0xFFFF_EA20	PIO Data Register 2
0xFFFF_EA24	PIO Direction Control Register 2
0xFFFF_EA28	PIO Effect Select Control Register 2

Table 12.2 PIO Register Address Assignments.

PIO Data Register 0..2 ('PioDataReg'0..2)

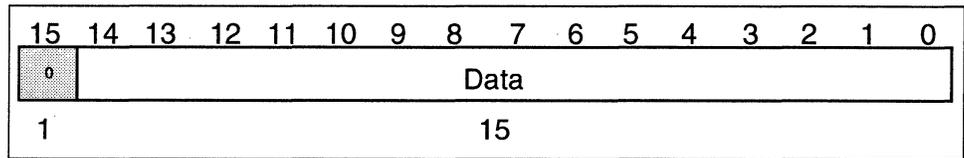


Figure 12.2 PIO Data Register ('PioDataReg').

PIO Data ('PIOData') Field:

Value	Action
'1'	PIO pin is high (default).
'0'	PIO pin is low.

Table 12.3 PIO Data ('PIOData') Field Encoding.

PIO Direction Register 0..2 ('PioDirReg'0..2)

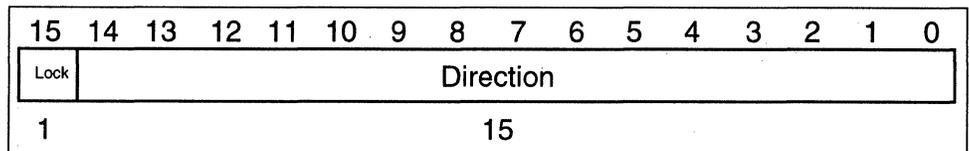


Figure 12.3 PIO Direction Register ('PioDirReg').

The PIO Direction Control Registers (see Figure 12.3) contain 16 bits each:

- The MSB is a lock bit (see Table 12.4).
- The bits in PIO Direction Register control whether the PIOs are inputs or outputs (see Table 12.5).

Lock ('Lock') Field:

Value	Action
'1'	Locks the Register from being altered by future writes.
'0'	No action (default).

Table 12.4 Lock ('Lock') Field Encoding.

Direction ('Dir') Field:

Value	Action
'1'	PIO pin is an output.
'0'	PIO pin is an input (default).
<p>Note: To avoid internal device damage, this field must be programmed carefully. When used in the input direction, any internal output or I/O driver—for example a serial port—must first be programmed to be tri-state or an input (default). Only then can the PIO 'Dir' field be safely changed from output to input.</p>	

Table 12.5 Direction ('Dir') Field Encoding.

PIO Effect Select Register 0..2 ('PioEffectSelReg'0..2)

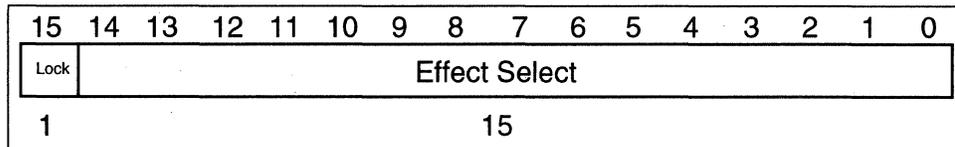


Figure 12.4 PIO Effect Select Register ('PioEffectSelReg').

The PIO Effect Select Registers (see Figure 12.4) contain 16 bits each:

- The MSB is a lock bit (see Table 12.6).
- The bits in PIO Effect Select Control Register 1 control whether the PIOs function as a special effect pin (e.g., Laser Printer, Serial Port, etc.) or as a general purpose PIO pin (see Table 12.7).

Lock ('Lock') Field:

Value	Action
'1'	Lock the Register from future writes.
'0'	No action (default).

Table 12.6 Lock ('Lock') Field Encoding.

Effect Select ('EffectSel') Field:

Value	Action
'1'	PIO pin is a special effect pin (default).
'0'	PIO pin is a general purpose pin.

Table 12.7 Effect Select ('EffectSel') Field Encoding.



Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. Many of the on-chip peripherals use and produce interrupts. Thus the R36100 includes an *Expansion Interrupt Controller*, as described in this chapter.

The Expansion Interrupt Controller, see Figure 13.1, works in conjunction with the CPO Status and Cause Registers. The Expansion Interrupt Controller steers the 20+ peripheral interrupts into one of the six CPO interrupt lines. In addition, since the DMA Controller can also use interrupts, each channel can select an interrupt from one of the peripherals.

This chapter provides an overview of the Expansion Interrupt programming interface and a complete description of the signal pins. Also included is an explanation on how expansion interrupts relate to typical internal and external systems.

Features

- Allows masking and status checking of all peripheral generated interrupts.
- Allows each DMA Controller Channel to receive an interrupt.

Block Diagrams

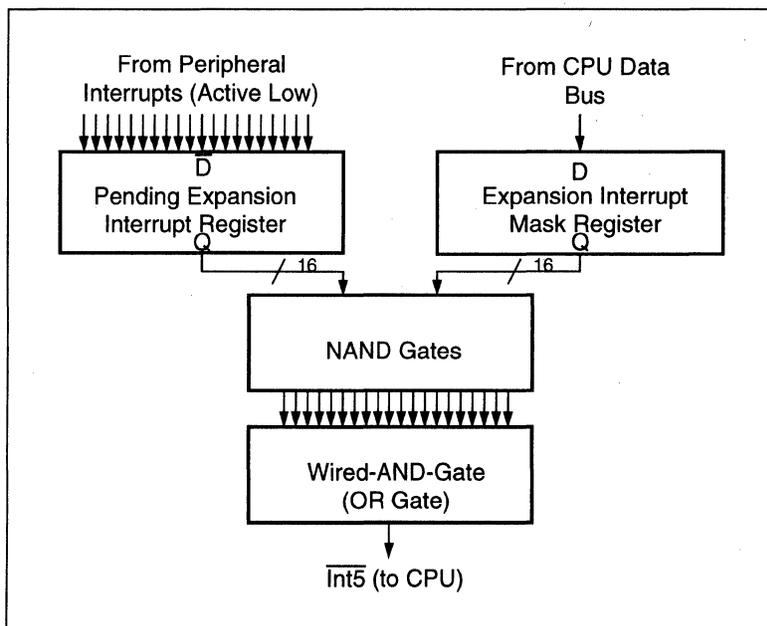


Figure 13.1 Expansion Interrupt Controller (to CPU Interrupt).

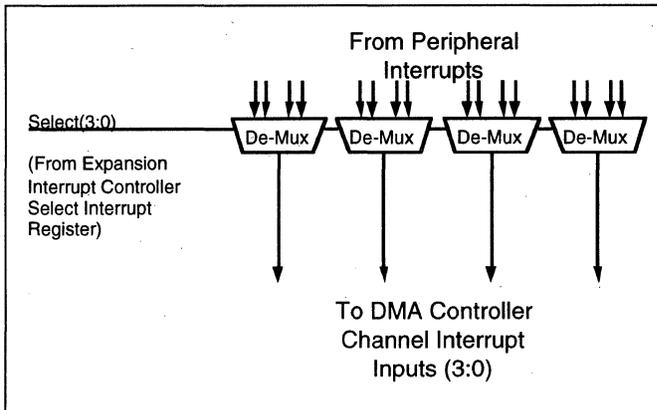


Figure 13.2 Expansion Interrupt Controller: Steering Interrupts to DMA Requests.

Overview

The Peripheral Expansion Interrupt Controller provides a means of steering the various 20+ peripheral generated interrupts (Figure 13.2). These peripheral interrupts are combined into a single CPU interrupt, Int(5). Each of the peripheral interrupts are stored (active high) in the Pending Expansion Interrupt Register on every system clock. If the corresponding mask bit in the Expansion Interrupt Mask Register is also set (active high), then the overall interrupt line Int(5) is set. At that point it is up to the CPU and ISR software to enable and handle Int(5).

The Peripheral Expansion Interrupt Controller also provides a means of steering a number of the peripheral interrupts to the DMA Channels. Each channel can select from a list of 4 peripheral interrupts. Table 13.1 provides an address map and description of the Expansion Interrupt Controller Register.

Pin Descriptions

Exception Signals

ExcSInt(2:0),

ExcInt(4:3)

Input

Processor Exception Synchronized Interrupt: These signals are functionally the same as the Int(4:0) signals of the R3000A. The synchronized interrupt inputs are internally synchronized by the R36100, and therefore may be generated by an asynchronous interrupt agent; the direct interrupts must be externally synchronized by the interrupt agent.

ExcSBrCond(3:2)

Input

Exception Synchronized Branch Condition Input: These input ports to the processor can use the Branch on Co-Processor Condition instructions to test their polarity. The branch condition inputs are synchronized by the R36100; therefore, they may be driven by an asynchronous source.

Phys. Address	Description
0xFFFF_EB00	Pend 0
0xFFFF_EB04	Mask 0
0xFFFF_EB10	Pend 1
0xFFFF_EB14	Mask 1
0xFFFF_EB86	DMA Select

Table 13.1 Expansion Interrupt Controller Register Address Assignments

Expansion Interrupt Mask Register 0..1
 ('ExpIntMaskReg0..1'),

and

Expansion Interrupt Pending Register 0..1
 ('ExpIntPendReg0..1'),

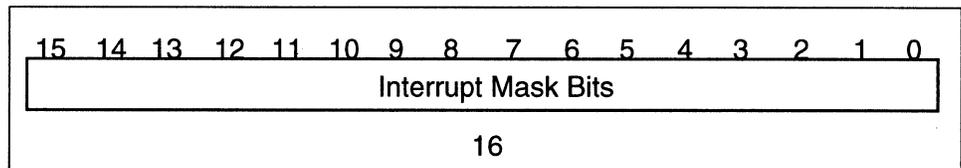


Figure 13.3 Expansion Interrupt Mask Register ('ExpIntMaskReg').

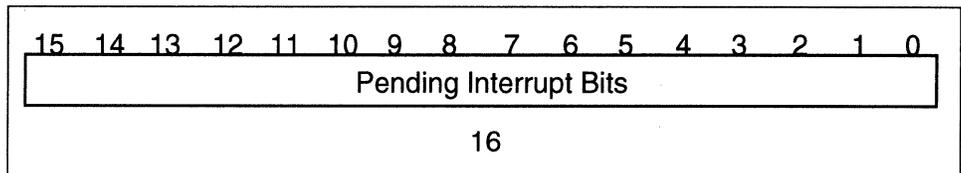


Figure 13.4 Expansion Interrupt Pending Register ('ExpIntPendReg').

A write to the Interrupt Pending Register (Figure 13.4) resets the register bit de-asserted low if the write data bit is a 1, or leaves the register bit with its present value if the write data bit is a 0. The pending interrupt register samples on every clock and holds an interrupt assertion until it is acknowledged by a write to the pending interrupt register. Note that the Expansion Interrupt Pending Register is different than the CPU CPO Pending Interrupt Field in that the Expansion Interrupt Pending Register has the added feature holding a pulsed interrupt until acknowledged. The Expansion Interrupt Mask Register is shown in Figure 13.3.

Additional programming instructions for the Expansion Interrupt Mask and Expansion Interrupts Pending registers are located in Table 13.2, Table 13.3, Table 13.4, and Table 13.5.

Reserved Low ('0') Field:

Must be written to '0' for future compatibility. Value when read is undefined.

Mask Bits ('Mask') Field:

Bit	Description
'4'	SerialInt
'3'	SerialRx_Req(1)
'2'	SerialRx_Req(0)
'1'	SerialTx_Req(1)
'0'	SerialTx_Req(0)
Note: Default values for the Pending and Mask Registers are all '0'.	

Table 13.2 Expansion Interrupt Mask Register 1 and Expansion Interrupt Pending Register 1 Bit Assignments.

Pending Bits ('Pend') Field:

Bit	Description
15	Reserved
14	CentReadInt
13	CentWriteInt
12	CentResetInt
11	LaserVideoBandInt
10	LaserFIFONotFull
9	LaserVideoPageInt
8	LaserVideoLineInt
7	TimerTC(2)
6	TimerTC(1)
5	TimerTC(0)
4	Reserved
3	DMADoneInterrupt3
2	DMADoneInterrupt2
1	DMADoneInterrupt1
0	DMADoneInterrupt0
Note: Default values for the Mask and Pending Registers are all '0'.	

Table 13.3 Expansion Interrupt Mask Register 0 and Expansion Interrupt Pending Register 0 Bit Assignments.

Value	Description
'1'	Interrupt pending.
'0'	Interrupt not pending.

Table 13.4 Pending Interrupt Field Encoding.

Value	Description
'1'	Interrupt enabled/allowed.
'0'	Interrupt disabled/disallowed.

Table 13.5 Interrupt Mask Field Encoding.

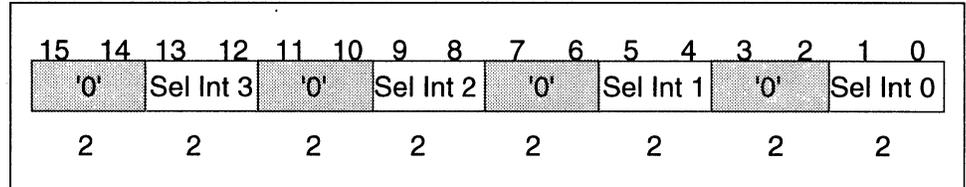
Expansion Interrupt DMA Select Register ('ExpIntDMASelReg')

Figure 13.5 Expansion Interrupt DMA Select Register ('ExpIntDMASelReg').

Select Interrupt 'SelInt()' Field:

Figure 13.5 gives the fields for the Expansion Interrupt DMA Select Register. The Select Interrupt field does a 1-in-4 select on the inputs to particular Interrupt Pending Register fields. The resulting input is passed to the Internal DMA Controllers. The input to the Interrupt Pending Register is used so that the DMA Controller does not need to acknowledge the Interrupt Pending Register to reset it, i.e., it ignores it. It is implied that the peripheral receiving or transmitting the data will de-assert its interrupt line when it receives a data strobe.

Additional programming instructions for this register are located in Table 13.6 and Table 13.7.

Note: Some interrupts are intentionally duplicated in multiple channels so that the system programmer can choose the relative priority of the interrupts.

Sel	IDMA(3)	IDMA(2)	IDMA(1)	IDMA(0)
3	CentWriteInt	CentReadInt	SerialTxReq(0)	LaserFIFOFull
2	SerialRx_Req(1)	SerialTx_Req(1)	SerialRx_Req(0)	SerialTx_Req(0)
1	TimerTc(2)	CentWriteInt	TimerTC(1)	TimerTC(0)
0	ExcInt(3)	ExcInt(2)	ExcInt(1)	ExcInt(0)

Table 13.6 DMA Channel versus Interrupt De-Multiplexer.

Value	Action
'15'-'4'	Reserved.
'3'	Select Interrupt 3.
'2'	Select Interrupt 2.
'1'	Select Interrupt 1.
'0'	Select Interrupt 0.

Table 13.7 Select Interrupt ('SelInt()') Field Encoding.



Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. The on-chip peripherals include *Timers* as described in this chapter.

This chapter will provide an overview of the Timer programming interface, a complete description of the signal pins and their waveforms, and discuss how the timers relate to typical internal and external systems. A block diagram of the R36100 Timers is located in Figure 14.1.

Features

- 3 16-bit Timer Channels with global 16-bit prescaler
- 3 TC/Gate pins
- Each Timer has:
 - 16-bit count register with selectable 16-bit frequency divider/prescaler of pipeline clock input
 - 16-bit compare register
 - TC control bit allowing auto reset vs. compare reg write ack for interrupts
 - Gate option control bit (gate option allows PWM counting or time stamping)
- Default Timer use is for a Real Time Clock/Timer.
- Timers have bus timeout control bit
 - Reset on bus start, gate on bus cycles
- Timer0 has 16-bit PWM low time compare register
 - Square Wave Generator

Block Diagram

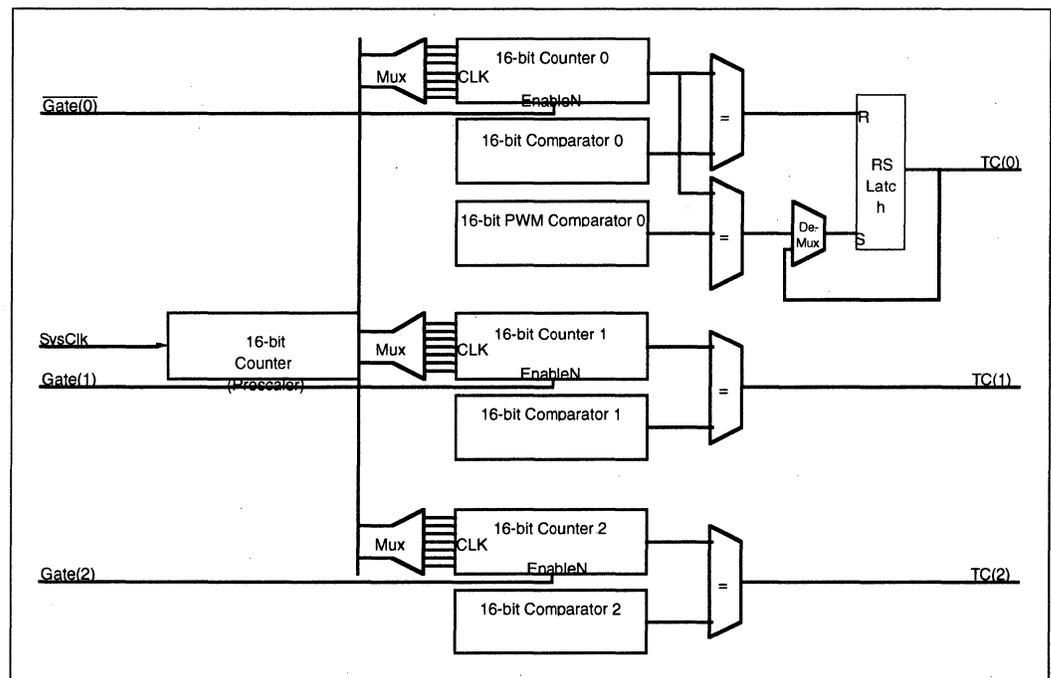


Figure 14.1 Block Diagram of the R36100 Timers.

Overview

The R36100 contains 3 timers. Each timer consists of a 16-bit count register as well as a 16-bit compare register. The count register resets to 0 and counts upward until it equals the compare register. When the count register equals the compare register, the \overline{TC} output is asserted and the count register is reset back to 0.

In order to expand the amount of time that the timers can handle, the timers use a common 16-bit prescaler counter. Each timer can be programmed to select a power-of-2 divisor of the prescaler as its fundamental base frequency for clock ticks. The prescaler counter itself is based off of the System Clock, \overline{SysClk} .

Using the default mode, each timer can be used as a real-time counter. Special effects include the following:

- Counter
- Real-time interrupt-based timer
- Bus timeout timer
- Gated clock external event counter

In addition to the above effects, timer channel 0 also has a PWM (Pulse Width Modulation, i.e., controllable duty cycle) compare register which controls the number of counter ticks that the timer output, \overline{TC} , remains asserted low, allowing for use as a PWM generator.

Pin Descriptions

Timer Peripheral Signals

TC(2:0),

Input/Output

TimerGate(2:0)

Gate: Active low input mode where the corresponding Timer() can increment its count whenever Gate() is asserted low. The Timer() stops counting whenever Gate() is high. Typically, the Timer() prescaler is set to the divide-by-1 mode so that each Gated clock corresponds to a Timer tick. Applications include Time Stamping, Pulse Width Measurement, and Timer expansion.

Terminal Count: Active low output mode is where the corresponding Timer() asserts \overline{TC} low whenever its Count Register equals its Compare Register. There is a one clock delay from Count equalling Compare until \overline{TC} is seen on the external pin due to internal synchronization. Normally \overline{TC} asserts low then immediately de-asserts back high after 1 Timer clock cycle [tbd. May be based on "tick"]. If the 'AckOnWrCompare' bit option in the Timer Control Register is asserted, then \overline{TC} de-asserts back high only when the Timer Compare Register is written. In this mode, the Timer can be used to generate an interrupt, and then the interrupt handler can acknowledge the interrupt. Timer 0 has a Pulse Width Modulation (PWM) Compare Register that can be programmed to de-assert \overline{TC} back high after a user programmed number of clock cycles. Applications include using the Timer to strobe external events, generating Real-time interrupts, Timer Expansion, and PWM control.

Register Descriptions

Table 14.1 provides a Timer Register Physical Address Map. Note that Big Endian software must offset these addresses by b'10 (0x2).

Address	Description
0xFFFF_E900 0xFFFF_E904	Timer Prescaler Count Register Reserved (must not read or write, used for internal PWM0 tcn low count.)
0xFFFF_E910 0xFFFF_E914 0xFFFF_E918 0xFFFF_E91C	Timer Count Register 0 Timer Compare Register 0 Timer PWM Register 0 Timer Control Register 0
0xFFFF_E920 0xFFFF_E924 0xFFFF_E92C	Timer Count Register 1 Timer Compare Register 1 Timer Control Register 1
0xFFFF_E930 0xFFFF_E934 0xFFFF_E93C	Timer Count Register 2 Timer Compare Register 2 Timer Control Register 2

Table 14.1 Timer Register Physical Address Map.

Timer Prescaler Count Register (‘TimerPrescalerCountReg’)

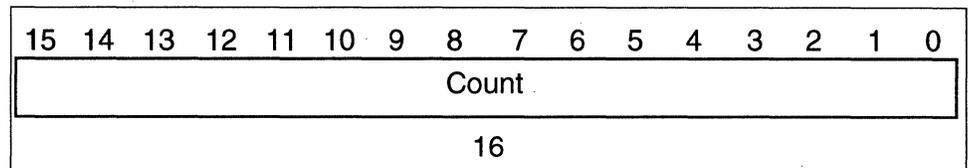


Figure 14.2 Timer Prescaler Count Register (‘TimerPrescalerCountReg’).

The prescaler counter starts at reset, and continuously counts in an upward direction, and wraps around on overflow. The reset default value is 0x0000. The prescaler counter uses the System Clock, $\overline{\text{SysClk}}$, as its fundamental base clock frequency. The Timer Prescaler Count Register is illustrated in Figure 14.2.

Timer Count Register 0..2 (‘TimerCountReg’0..2)

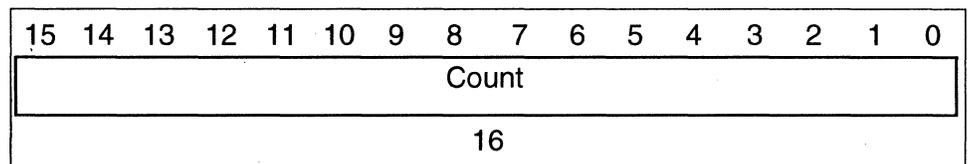


Figure 14.3 Timer Count Register (‘TimerCountReg’).

The Count Register, shown in Figure 14.3, is a 16-bit register that contains the present count value. The Count Register increments on every prescaled clock tick. The default value at reset for the Count Register is 0x0000.

The Count Register does not count if $\overline{\text{PIOInputGate}}$ is enabled via the Timer Control Register and the input gate signal, $\overline{\text{Gate}}$ is high. The Count Register is 16-bit Readable and Writable if $\overline{\text{LockCountAndCompare}}$ control bit is not active. The Count Register ignores the $\overline{\text{Gate}}$ input if the $\overline{\text{Gate}}$ input control bit is not turned on.

**Timer Compare Register 0..2
(‘TimerCompareReg’0..2)**

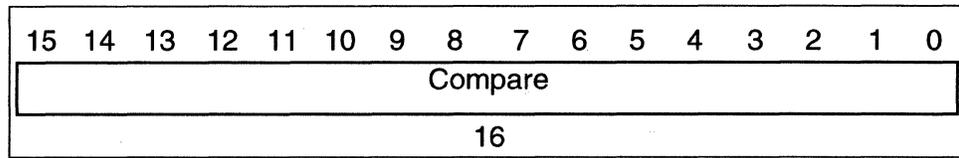


Figure 14.4 Timer Compare Register (‘TimerCompareReg’).

The Compare Register, illustrated in Figure 14.4, is a 16-bit register containing the value that will reset the counter when the Count Register is equal to it. The default value at reset for the Compare Register is 0xFFFF. This register is 16-bit readable and writable if the 'Lock-CountAndCompare' control bit is not active. If written and 'WriteCompareAck' Timer Control Register bit is active, then it brings \overline{TC} back to inactive high.

**Timer Pulse Width Modulation Register 0
(‘TimerPWMReg’0)**

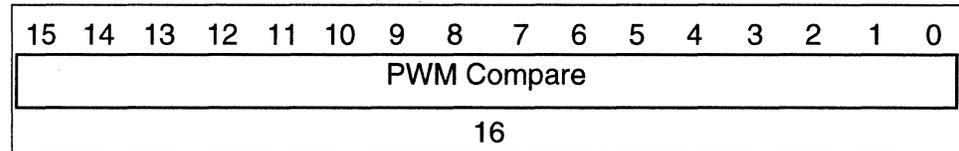


Figure 14.5 Timer Pulse Width Modulation Register (‘TimerPWMReg’).

The Timer Pulse Width Modulation Register, illustrated in Figure 14.5, has a 16-bit value which brings PWM \overline{TC} output back high 'N' prescaled clocks after \overline{TC} goes low. Reset default value is 0x0000.

By programming various values for the compare and the PWM compare registers, the duty cycle of the \overline{TC} output can be varied. For instance, by using a compare value equal to a PWM-1 value, the duty cycle will be 50/50. For the \overline{TC} output to always be high, the \overline{TC} PIO pin should be temporarily programmed as a general purpose output with a value of '1'. Additional programming information for this register is located in Table 14.2.

Value	Action
0	high after 1 clock
1	high after 2 clocks
compare	high for 1 clock
> compare	always low

Table 14.2 Timer Pulse Width Modulation Register (‘TimerPWMReg’) Bit Fields.

**Timer Control Register 0..2
(‘TimerControlReg’0..2)**

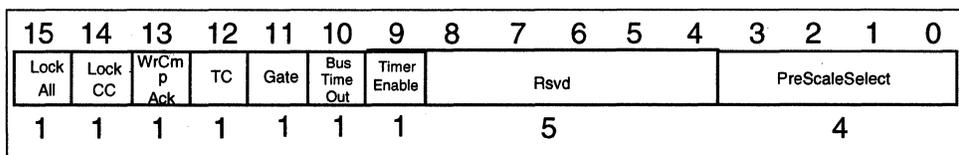


Figure 14.6 Timer Control Register (‘TimerControlReg’).

Figure 14.6 illustrates the Timer Control Register. Additional programming information is located in Table 14.3, Table 14.4, Table 14.5, Table 14.6, Table 14.7, Table 14.8, Table 14.9, Table 14.10, and Table 14.11.

Bit	Description
15	Lock
14	LockCountAndCompare
13	AutoAck vs WriteCompareAck
12	PIO is Output \overline{TC}
11	PIO is Input \overline{Gate}
10	BusTimeout
9	TimerEnable
8:4	Reserved. Must be written as '0'
3:0	Prescale 1 of 16 Select

Table 14.3 Timer Control Register ('TimerControlReg') Bit Assignments.

Lock ('Lock') Field:

Value	Action
'1'	Control locked from future writes.
'0'	No action (default).

Table 14.4 Lock ('Lock') Field Encoding.

Lock Count and Compare ('LockCC') Field:

Value	Action
'1'	Count and Compare Registers locked from future writes.
'0'	No action (default).

Table 14.5 Lock Count and Compare ('LockCC') Field Encoding.

Write Compare Ack ('Ack') Field:

Value	Action
'1'	WriteCompareAck requires a write to the compare register to set \overline{TC} back high.
'0'	\overline{TC} back high 1 clock after \overline{TC} asserts low. (default)

Table 14.6 Write Compare Ack ('Ack') Field Encoding.

PIO is Output \overline{TC} ('TC') Field:

The PIO is Output \overline{TC} field must be identically programmed to the actual PIO pin.

Value	Action
'1'	If PIO pin is programmed elsewhere as an output, then \overline{TC} is driven out of that PIO output.
'0'	Disable external \overline{TC} . (default)

Table 14.7 PIO is Output \overline{TC} ('TC') Field Encoding.

PIO is Input \overline{Gate} ('Gate') Field:

The PIO is Input Gate field must be programmed identically to the actual PIO pin.

Value	Action
'1'	If PIO pin is programmed elsewhere as an input, then \overline{TC} is driven in from that PIO input.
'0'	Does not enable external \overline{TC} . (default)

Table 14.8 PIO is Input Gate ('Gate') Field Encoding.

BusTimeout ('BTO') Field:

Value	Action
'1'	Enable Bus Timeout feature. Reset counter to 0 on the beginning of each external bus cycle.
'0'	Disable use Bus Timeout feature. (default)

Table 14.9 BusTimeout ('BTO') Field Encoding.

Timer Enable ('TimerEn') Field:

This field enables the timer to count.

Value	Action
'1'	Stop
'0'	Timer Enabled (default)

Table 14.10 Timer Enable ('TimerEn') Field Encoding

Prescaler Select ('PSel') Field:

Note that changing the Prescaler Select can take up to $2 \cdot 16$ clocks to take effect due to internal synchronization.

Value	Prescaler Divide-By Frequency
'F'	div 2^{15} (divide by 32768)
...	...
'2'	div 2^2 (divide by 4)
'1'	div 2^1 (divide by 2)
'0'	div 2^0 (divide by 1) (default)

Table 14.11 Prescaler Select ('PSel') Field Encoding.



Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. The R36100 contains two serial port channels. One of the many on-chip peripherals is the *Serial Port* interface (see Figure 15.1), as described in this chapter.

Chapter 15 provides an overview of the Serial Port register interface, a description of the signal pins as well as various aspects of the signal timing. Figure 15.1 is a serial channel interface block diagram.

Features

- 2 serial channels
- Asynchronous mode features such as:
 - Bits per char, stop bit, parity, and error detection options
 - Modem handshaking signals
- Synchronous byte-mode features such as:
 - Monosync/Bisync/SDLC/HDLC packet frame options
 - Sync, flag, address field, abort, CRC generation and detection
 - 10 deep Frame Status FIFO
- 3 deep Receiver Data FIFO
- 2 deep Transmitter Data FIFO
- Data Encoder/Decoder (ENDEC) with NRZ, NRZI, FM0, and FM1 options
- Phase lock loop (PLL) for synchronous clock regeneration
- Baud Rate Clock Generator
- High Speed data rates (up to 1/4 speed of the CPU SysClk)

Block Diagram

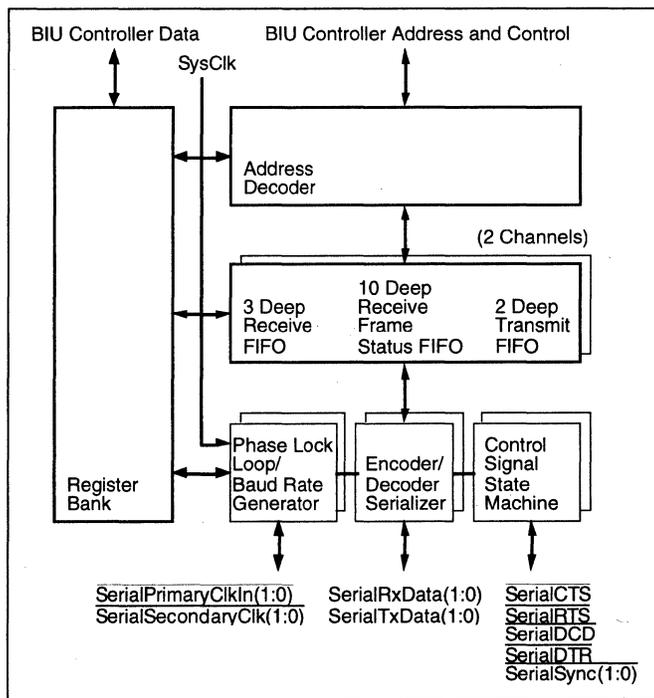


Figure 15.1 Serial Channel Interface Block Diagram.

Overview

The R36100 contains two independent serial port channels. Either port can be used in asynchronous mode to support RS-232 protocols or in synchronous mode to support Monosync, Bisync, SDLC, and HDLC protocols.

In the asynchronous mode, the Serial Ports are capable of full-duplex communication either via CPU polling or interrupt service routines. In addition, the internal DMA Controller may optionally be used to automate message buffer transfers from main memory to/from the Serial Ports. Various asynchronous features include:

- character length, parity, and stop bit options
- Input clocking options include:
 - using an external oscillator directly
 - using an external oscillator or using the internal CPU provided SysClk via the internal Baud Rate Clock Generator of each channel

In the typical asynchronous 16 clock sampling mode, the serial ports can support up to 1Mbaud (with 16MHz CPU, 2Mbaud with 33MHz CPU). On errors, such as parity or framing errors, the serial ports can notify the CPU either via interrupt or via CPU polling. Modem handshaking signals including CTS/RTS, DTR, and DCD can optionally be controlled. A 3 deep receive data FIFO and a 2 deep transmit data FIFO are provided to buffer data.

In the typical synchronous 16 FM PLL clock sampling mode, the serial ports can support 300-1Mbaud (with 16MHz CPU, 2Mbaud with 33MHz CPU). Non-PLL modes can support up to 4Mbaud with the 16MHz CPU (8.25Mbaud with the 33MHz CPU). A 3 deep receive data FIFO and a 2 deep transmit data FIFO are provided to buffer data.

While in the synchronous mode, the Serial Ports are capable of full-duplex communication either via CPU polling, interrupt service routines, or semi-autonomously via the internal DMA Controller. Various features include:

- Sync, 0-bit, CRC, Flag, and abort character generation/insertion and detection/checking/deletion
- NRZ, NRZI, and FM encoder/decoder (ENDEC)
- 10-deep Frame Status FIFO

Input clocking options for each channel include: use of an external oscillator, use of the on-chip PLL, and use of the channel's internal Baud Rate Clock Generator as clocked via the CPU's SysClk or the external oscillator.

DMA Connections

The Serial Port Controller is connected internally to the DMA controller in full duplex mode. For example, while in the full duplex mode, DMA channel 0 is connected by the programmer to the Serial Channel 0 transmitter, DMA channel 2 to the serial channel 0 receiver, DMA channel 1 to the serial channel 1 transmitter, and DMA channel 3 to the serial channel 1 receiver.

By assigning a separate DMA Channel to each data stream, the internal DMA Controller can manage the details of moving the individual datum to or from the Serial Ports from or to main memory data stream buffers while the CPU can manage and maintain the higher level message or frame structures.

DMA Data Movement

The DMA controller can access the Serial Port interface in the same manner as the CPU and in the same number of clock cycles.

The Serial Port is connected to InternalData(7:0) of the internal data bus, regardless of the endianness of the system. It only supports byte data transfers (sb or lb instructions) and will ignore the other 3 data lanes. The port width can be programmed as 8, 16, or 32-bits; however, (mini-) burst transfers are not supported, so the 32-bit port width is recommended. When the CPU is in the Big Endian byte ordering mode, the DMA controller must be programmed to appropriately funnel the data to and from the Serial Port to the InternalData(7:0) byte lane by using an address offset of 'b11.

Serial Data Speed Calculations

The analysis presented here estimates the maximum data rate that this device can sustain using DMA techniques. The calculations are based on the following assumptions:

- DMA access to the Serial Port (read or write) = $5 \overline{\text{SysClk}}$'s.
- DMA access to the DRAM system (read or write) = $3 \overline{\text{SysClk}}$'s (actual number is system dependent).
- DMA bus arbitration = $4 \overline{\text{SysClk}}$'s (actual average is system dependent).
- DMA bus release to the CPU = $1 \overline{\text{SysClk}}$ (actual is 0-2).

It is further assumed that every DMA operation will consist of:

- DMA access to Serial Port + DMA access to DRAM + DMA bus release to CPU = $5 + 3 + 1 = 9 \overline{\text{SysClk}}$'s.
- The DMA bus arbitration is added only once in the calculation, because it should usually occur in parallel with other operations and should not constitute a major portion of the transfer.
- DMA should support a minimum of 2.048 Mbit/sec full duplex synchronous transfer on every channel to support maximum Serial Port rates:
 - Data rate (bit) = 2.048Mbit/sec
 - Data rate (byte) = (2.048Mbit/sec) / 8 = 0.256MByte/sec
 - Time per byte = $1 / (0.256\text{MBytes/sec}) = 3.90625\mu\text{sec/byte}$
 - Number of clock cycles (for 33MHz) = $3.90625\mu\text{sec} / 30\text{nsec} = 130 \overline{\text{SysClk}}$ cycles

For a full duplex operation, 2 bytes need to be transferred between the Serial port and the DRAM per channel (transmit and receive) during this time frame. (The actual requirement is slightly more lenient in that 6 bytes of receive data must occur every 3 of these time frames). For 2 full duplex channels, the actual number of bytes to be transferred is then 4 bytes.

- Number of clock for byte transfer = $(9 \overline{\text{SysClk}}) * 4 = 36 \overline{\text{SysClk}}$'s
- The total number (including DMA bus arbitration) = $36 + 4 = 40 \overline{\text{SysClk}}$'s
- The bus utilization of the Serial Port DMA operation = $40/130 = 31\%$.

This number is worst case estimate; however, it doesn't include the clock cycles required to set up the DMA controller registers for the CPU to respond to any Serial Port interrupt.

Compatibility Issues

Although the R36100 Serial Ports are similar in data link layer functionality and contain many of the same resources as other serial port devices, the R36100 Serial Ports are not specifically designed to be 100% functionally compatible with any other specific serial port device. Areas of difference include pin AC/DC electrical and timing characteristics, register assignments, bit assignments, and detailed functional specification.

In many communication situations, both software and hardware use an application to transport data link/physical layer approach. The middle level/layer protocols, such as LocalTalk, can easily accommodate the R36100 Serial Ports, with the addition of an assembly/low level C-code device driver specifically written/ported for the R36100.

Pin Definitions

Serial Port Signals

Serial port pins that are input only or output only are multiplexed with PIO registers such that they can be reprogrammed to be general purpose inputs or outputs. The exceptions are `SerialRxData(0)`, `SerialTxData(0)` and `SerialPrimaryClk(0)` which are always dedicated to their serial port function.

`SerialRxData(1:0)` **Input**

Serial Port Receiver Data: The Serial Port receives the serial input data stream via this active high input. For typical RS232-C serial connections, an external transceiver inverts and level shifts this signal from +/-12V.

`SerialTxData(1:0)` **Output**

Serial Port Transmitter Data: The Serial Port sends the serial transmitter output data stream via this active high output. For typical RS232-C serial connections, an external transceiver inverts and level shifts this signal to +/-12V.

`SerialPrimaryClkIn(1:0)` **Input**

Serial Port Primary Receiver/Transmitter Clock Input Negated: Optionally provides the receiver clock, the transmitter clock; or the base clock for the Baud Rate Clock Generator or the PLL clock regenerator. The functionality of the pin is selected by programming the internal registers of the serial port. In addition, the `SerialPrimaryClkIn()` pins can be programmed to be substituted with `SysClk` when used as the baud-rate clock generator input clock.

If the receiver clock is provided by `SerialPrimaryClkIn()`, the serial data stream on the `SerialRxData()` will be sampled on the de-asserting edge of the receiver clock in NRZ and NRZI modes. In FM modes, on the other hand, the input serial data stream on `SerialRxData()` will be sampled on both the asserting and de-asserting clock edges.

`SerialSecondaryClk(1:0)` **Input/Output**

Serial Port Secondary Receiver/Transmitter Clock Negated Input: When this pin is programmed to be an active low input, it can optionally provide the receiver clock or the transmitter clock (bypassing the Baud Rate Clock Generator). If the transmitter clock is provided through `SerialSecondaryClk()`, the serial data stream on the `SerialTxData()` will be sampled on the de-asserting edge of the transmitter clock in NRZ and NRZI modes. In FM modes, on the other hand, the output serial data stream on the `SerialTxData()` will be sampled on both the asserting and de-asserting clock edges.

Serial Port Secondary Receiver/Transceiver Clock Negated Output: When this pin is programmed to be an active low output, it can provide the transmitter clock output, Baud Rate Clock Generator input, or the PLL clock regenerator input.

SerialSync(1:0)**Input/Output****Serial Port Synchronization Negated:**

Asynchronous Mode Input: active low input which has no effect other than to change the state of the Sync status register bit. This input can be used as the Ring Indicator modem handshake signal.

External Synchronization Mode Input: Active low input which must be asserted for two clock cycles after the sync character is received in order to initiate the beginning of a frame.

Internal Synchronization Mode Output: Active low output asserts whenever a sync/flag character is recognized.

SerialCTS(1:0)**Input**

Serial Port Clear to Send Negated: Active low input indicating that a communications agent (such as modem) is ready to send data. When programmed in the Auto-Handshaking mode, the Serial Port transmitter can be controlled (enabled) by this input.

SerialRTS(1:0)**Output**

Serial Port Request to Send Negated: Active low output indicating that the Serial Port has data to transmit to a communications agent (such as a modem). When programmed to the asynchronous Auto-Handshaking mode, SerialRTS asserts whenever the transmitter is not empty.

SerialDCD(1:0)**Input**

Serial Port Data Carrier Detect Negated: Active low input indicating that a communications agent is receiving valid data from the communication line. When the channel is programmed to be in the Auto-Handshaking mode, the Serial Port receiver can be controlled (enabled) by this input.

SerialDTR(1:0)**Output**

Serial Port Data Terminal Ready Negated: This active low output can be used as a general purpose output pin by writing to the DTR status register bit.

Register Definitions

The R36100 Serial Ports use a concept called "MetaRegisters." The MetaRegisters provide a port like interface into the actual internal control registers of each channel. In general, the MetaRegister must first be written to with a Pointer Value indicating the internal register number that is to be accessed next (see Table 15. 2, Table 15. 3, Table 15. 4, and Table 15. 5 for more detailed programming information). Next, the read or write, from the channel's MetaRegister, will access that channel's internal register; therefore, internal registers are accessed with the following paired command sequence:

1. internal register read: write followed by read to the MetaRegister
2. internal write: write followed by write to the MetaRegister

There are a few internal registers directly accessed by the CPU to speed up their access--such as the receive data or transmit data FIFOs--which can be accessed by reading or writing the Data Register, respectively.

The MetaRegisters use byte addressing; thus, big endian systems must use an offset of b'11 (0x3). Note that because of the pointer before read/write command register hardware construct, it is necessary for software to guarantee that no interrupts that may context switch to a Serial Port Command register related ISR occur between a register pointer/command pair. In any case, reads and writes always use InternalData(7:0) and ignore InternalData(31:8).

Phys. Address	Description
0xFFFF_E800	Serial Port MetaRegister for Channel 1
0xFFFF_E804	Serial Port Data Register for Channel 1
0xFFFF_E808	Serial Port MetaRegister for Channel 0
0xFFFF_E80C	Serial Port Data Register for Channel 0
0xFFFF_E810 (read)	Serial Port Reset Register for Both Channels

Table 15. 1 List of Serial Port MetaRegisters.

Pointer Address	Description
0x00	LSB Status Register for each Channel
0x01	MSB Status Register for each Channel
0x02	Reserved when reading Channel 0, Interrupt Vector when reading Channel 1
0x03	Interrupt Pending Register for each Channel
0x08	Receiver Data FIFO for each Channel
0x0A	Clock Status Register for each Channel
0x0C	LSB Baud Rate Generator Compare Register for each Channel
0x0D	MSB Baud Rate Generator Compare Register for each Channel
0x0F	MSB Interrupt Enable Register for each Channel

Table 15. 2 Serial Port Internal Read Registers for each Channel (as pointed to by the MetaRegister).

Pointer Address	Description
0x00	Reset Register for each Channel
0x01	LSB Interrupt Enable Register 0 for each Channel
0x02	Reserved Low ('0') when writing Channel 0, Interrupt Vector when writing Channel 1
0x03	Receiver Control Register for each Channel
0x04	LSB Mode Register for each Channel
0x05	Transmitter Control Register for each Channel
0x06	LSB Sync Char Register for each Channel
0x07	MSB Sync Char Register for each Channel
0x08	Transmitter Data FIFO Register for each Channel
0x09	Reset Register for each Channel
0x0A	MSB Mode Register for each Channel
0x0B	Clock Control Register for each Channel
0x0C	LSB Baud Rate Generator Compare Register for each Channel
0x0D	MSB Baud Rate Generator Compare Register for each Channel
0x0E	PLL Control Register for each Channel
0x0F	MSB Interrupt Enable Register for each Channel

Table 15. 3 Serial Port Internal Write Registers for each Channel (as pointed to by the Meta Register).

Pointer Address	Description
0x06	LSB Frame Status Register for each Channel (read)
0x07	MSB Frame Status Register for each Channel (read)

Table 15. 4 Serial Port Internal Read Registers for each Channel as pointed to by the MetaPtrRegister 0xF.

Pointer Address	Description
0x07	Meta Meta Control Register for each Channel (write)

Table 15. 5 Serial Port Internal Write Register for each Channel as pointed to by the MetaPtrRegister 0x0F.

Data Register Usage

On a read, the receive data register is read and the receive data FIFO entry is removed. On a write, the transmit data register is written. The transmit FIFO must be non-full.

MetaRegister Usage

This register is used as a paired access. The first access is always a write which is used to write a 3-bit pointer that points to the internal register that is to be read or written. The next read or write is the actual command read/write, after which the pointer automatically returns to 0.

Interrupt service routines should take care to either not use the serial port directly, or, if it does, to read the pointer value first and return it to the same pointer value when completed. Typically, an ISR uses a semaphore to note if a character needs to be read/written and then passes this semaphore to a higher level routine to do the actual Serial Port data read/write.

MetaMetaRegister Usage

By enabling the Frame Status FIFO via bit 2 of the Interrupt Enable Write Register, write register 7 accesses a different function than its MetaFunction as MSB Sync Char Register to a new function, called Meta Meta Register. In addition read registers 7:6 become accessible as the MSB and LSB Frame Status FIFO read registers.

Figure 15.2 gives the fields for the Serial LSB Status Register, with bit assignments and field encodings listed in Table 15.6 and Table 15.7. The Serial MSB register fields are shown in Figure 15.3, with field encodings listed in Table 15.8 and Table 15.9. And the Serial Interrupt Vector is shown in Figure 15.4, with bit assignments and field encodings for this register listed in Table 15.10 and Table 15.11 on page 11.

Serial LSB Status Register ('SerialLSBStatusReg(1:0)') (Read Pointer 0x00)

7	6	5	4	3	2	1	0
Break	Underrun	CTSN	Sync	DCDN	TxFIFO Empty	Count	Rx NotEmpty
1	1	1	1	1	1	1	1

Figure 15.2 Serial LSB Status Register ('SerialLSBStatusReg').

Bit	Assignment
7	Break (asynchronous mode) or Abort (SDLC Mode) ('Break')
6	Transmitter Underrun ('Underrun')
5	CTS Status ('CTSN')
4	Sync Status ('Sync')
3	DCD Status ('DCDN')
2	Transmitter FIFO Empty Status ('TxFIFOEmpty')
1	Count Rollover Status ('Count')
0	Receiver FIFO Not Empty Status ('RxNotEmpty')

Table 15.6 Serial LSB Status Register ('SerialLSBStatusReg') Bit Assignments.

LSB Status Bits

Value	Status
'1'	Status Condition is true
'0'	Status Condition is false

Table 15.7 LSB Status Bit Field Encoding.

**Serial MSB Status Register ('SerialMSBStatusReg(1:0)')
(Read Pointer 0x01)**

7	6	5	4	3	2	1	0
ValidSDLC Frame	FrameErr	RxOverrun	ParityErr	Rsvd	Rsvd	Rsvd	TxDone
1	1	1	1	1	1	1	1

Figure 15.3 Serial MSB Status Register ('SerialMSBStatusReg')

Bit	Assignment
7	SDLC Valid End of Frame Status ('ValidSDLCFrame')
6	CRC/Frame Error Status ('FrameErr')
5	Receiver Overrun Error Status ('RxOverrun')
4	Parity Error Status ('ParityErr')
3:1	Reserved ('Rsvd')
0	Asynchronous Transmitter Done with Char Status ('TxDone')

Table 15. 8 Serial MSB Status Register ('SerialMSBStatusReg') Field Encoding.

Value	Status
'1'	Status Condition is true
'0'	Status Condition is false

Table 15. 9 MSB Status Bit Field Encoding.

Serial Interrupt Vector ('SerialIntVector(1)') (Pointer 0x02)

7	6	5	4	3	2	1	0
0					IntVector		
5					3		

Figure 15.4 Serial Interrupt Vector ('SerialIntVector(1)').

Bit	Assignment
7:3	Reserved Low ('0').
2:0	Interrupt Vector ('IntVector')

Table 15. 10 Serial Interrupt Vector Register ('SerialRxControlReg') Bit Assignments.

Value	Action
'7'	Exception Channel 0
'6'	Receiver FIFO Not Empty Channel 0
'5'	External Status Channel 0
'4'	Transmitter FIFO Empty Channel 0
'3'	Exception Channel 0
'2'	Receiver FIFO Not Empty Channel 0
'1'	External Status Channel 0
'0'	Transmitter FIFO Empty Channel 0

Table 15. 11 Interrupt Vector ('IntVector') Field Encoding.

Serial LSB Sync Character/Address Register
 ('SerialLSBSyncCharReg(1:0)')
 (Pointer 0x06),

Serial MSB Sync Character/Address Register
 ('SerialMSBSyncCharReg(1:0)')
 (Pointer 0x07)

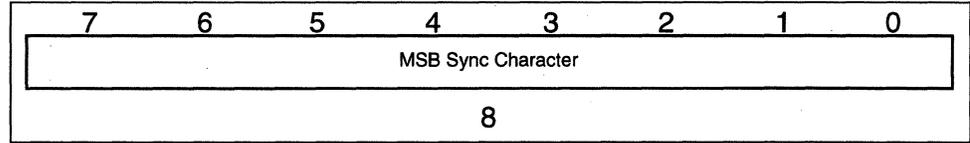


Figure 15.5 Serial MSB Sync Character/Address Register ('SerialMSBSyncCharReg').

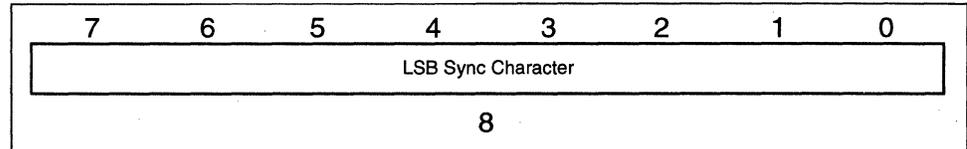


Figure 15.6 Serial LSB Sync Character/Address Register ('SerialLSBSyncCharReg').

For additional programming information, please see the Write Pointer description of the Serial Sync Character/Address Register in Figure 15.5 and Figure 15.6.

Receiver Data FIFO Register
 ('SerialDataFIFOReg(1:0)') (Read Pointer 0x08)

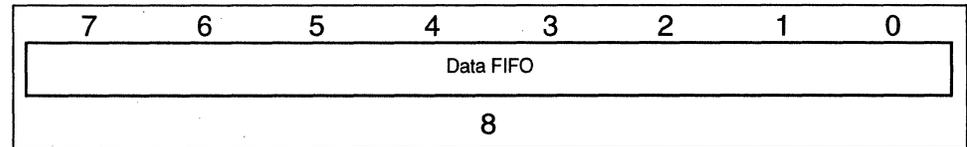


Figure 15.7 Receiver Data FIFO Register ('SerialDataFIFOReg').

The SerialDataFIFOReg is also directly accessible by the CPU by reading address 0xFFFF_E800 for Channel 0 or 0xFFFF_E808 for Channel 1. This register is shown in Figure 15.7.

Serial Clock Status Register
 ('SerialClkStatusReg(1:0)') (Read Pointer 0x0A)

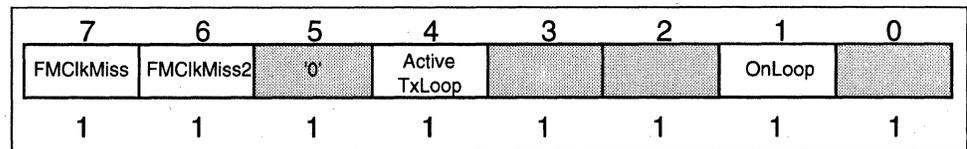


Figure 15.8 Serial Clock Status Register ('SerialClkStatusReg').

Bit	Assignment
7	FM Clock Missing ('FMClkMiss')
6	Two FM Clocks Missing ('FMClkMiss2')
5	Reserved Low ('0')
4	Active Transmitter On Loop ('ActiveTxLoop')
3:2	Reserved Low ('0')
1	On Loop Status ('On Loop')
0	Reserved Low ('0')

Table 15. 12 Serial Clock Status Register ('SerialClkStatusReg') Bit Assignments.

Value	Status
'1'	Status Condition is true
'0'	Status Condition is false

Table 15. 13 Clock Status Bit Field Encoding.

The Serial Clock Status register fields are shown in Figure 15.8, with bit assignments and field encodings for this register listed in Table 15. 12 and Table 15. 13.

Serial LSB Baud Rate Generator Compare Register ('SerialMSBBaudRateGenCompareReg(1:0)' (Pointer 0x0C)

Serial LSB Baud Rate Generator Compare Register ('SerialLSBBaudRateGenCompareReg(1:0)' (Pointer 0x0D)

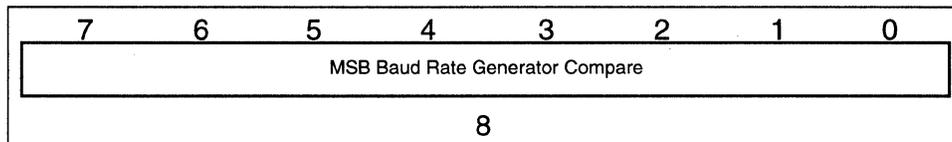


Figure 15.9 Serial MSB Baud Rate Clock Generator Compare Register ('SerialMSBBaud-CompareReg').

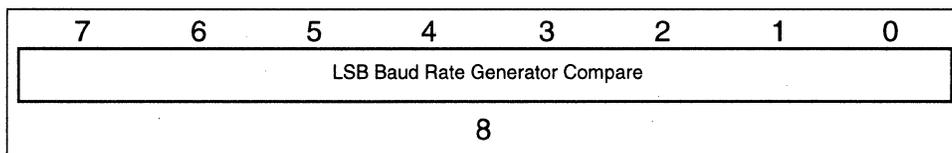


Figure 15.10 Serial LSB Baud Rate Clock Generator Compare Register ('SerialLSBBaud-CompareReg').

For additional programming information, please see the Write Pointer description for the Serial Baud Rate Clock Generator Compare Register in Figure 15.9 and Figure 15.10.

**Serial MSB Interrupt Enable Register ('SerialIntEnReg(1:0)') (Pointer 0x0F),
Meta Pointer Register ('MetaPointerReg(1:0)') (Pointer 0x0F)**

7	6	5	4	3	2	1	0
BreakIntEn	Underrun IntEn	CTSIntN	SyncIntEn	DCDNIntEn	FrameFIFO En	ZeroIntEn	MetaMeta RegEn
1	1	1	1	1	1	1	1

Figure 15.11 Serial MSB Interrupt Enable Register ('SerialIntEnReg').

Bit	Assignment
'7'	Break or Abort Interrupt Enable ('BreakIntEn')
'6'	Transmitter Underrun or EOM Condition Interrupt Enable ('UnderrunIntEn')
'5'	CTS Interrupt Enable ('CTSIntEn')
'4'	Sync Char Interrupt Enable ('SyncIntEn')
'3'	DCD Interrupt Enable ('DCDIntEn')
'2'	Frame Status FIFO Enable (MetaMetaRegister B) Enable ('FrameFIFOEn')
'1'	Baud Rate Clock Generator Zero Condition Interrupt Enable ('ZeroIntEn')
'0'	MetaMeta Control Register Enable ('MetaMetaRegEn')

Table 15.14 Serial MSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments.

For additional programming information, refer to the Write Pointer description for the Serial MSB Interrupt Enable Register in Figure 15.11. Bit assignments and field encodings for this register are listed in Table 15.14.

**Serial Meta Register ('SerialMetaRegister(1:0)')
(Write Pointer 0x00)**

7	6	5	4	3	2	1	0
ResetA		ResetB / Pointer					
2		6					

Figure 15.12 Serial Meta Register ('SerialMetaRegister').

Pointers and Reset B Commands must be written on separate writes.

Bit	Assignment
7:6	Reset A Command
5:0	Reset B Command / Pointer

Table 15.15 Serial Meta Register ('SerialMetaRegister') Bit Assignments.

If the Reset Transmitter Underrun Condition has been given, then the CRC is inserted (or SDLC Abort/Flag if enabled) at the end of a transmitter underrun.

Value	Action
'3'	Reset Transmitter Underrun Condition
'2'	Reset Transmitter CRC Checker
'1'	Reset Receiver CRC Generator
'0'	No Effect.

Table 15. 16 Reset A Command ('ResetA') Field Encoding.

Value	Action
'38'	Reset Highest Priority Interrupt
'30'	Reset Exception Condition
'28'	Reset Transmitter Pending Interrupt
'20'	Enable Interrupt on Next Receiver Character
'18'	Send Abort
'10'	Reset External Status Interrupts

Table 15. 17 Reset B Command ('ResetB') Field Encoding.

Value	Action
'b001111'	Internal Register F
'b001110'	Internal Register E
...	...
'b000000'	Internal Register 0

Table 15. 18 Pointer Value ('Pointer') Field Encoding.

The Serial Meta Register is shown in Figure 15.12. Refer to Table 15. 15, Table 15. 16, Table 15. 17, and Table 15. 18 for additional programming information.

Serial LSB Interrupt Enable Register ('SerialIntEnReg(1:0)')
 (Write Pointer 0x01)

7	6	5	4	3	2	1	0
0	0	0	RxIntEn	Parity IntEn	TxIntEn	ExtStatus IntEn	
1	1	1	2	1	1	1	

Figure 15.13 Serial LSB Interrupt Enable Register ('SerialIntEnReg').

Bit	Assignment
7	Reserved High ('1')
6	Reserved High ('1')
5	Reserved High ('1')
4:3	Receiver Interrupt Enable ('RxIntEn')
2	Parity Interrupt Enable ('ParityIntEn')
1	Transmitter Interrupt Enable ('TxIntEn')
0	External Status Interrupt Enable ('ExtStatusIntEn')

Table 15. 19 Serial LSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments.

Receiver Interrupt Enable Mode ('RxIntEn')

Value	Action
'3'	Enable Receiver Interrupt Only on Exception
'2'	Enable Receiver Interrupt on Any Character or Exception
'1'	Enable Receiver Interrupt only on First Character or Any Exception
'0'	Disable Receiver Interrupt

Table 15. 20 Receiver Interrupt Enable Mode ('RxIntEn') Field Encoding.

Parity Interrupt Enable Mode ('ParityIntEn')

Value	Action
'1'	Enable Parity Interrupt
'0'	Disable Parity Interrupt

Table 15. 21 Parity Interrupt Enable ('ParityIntEn') Field Encoding.

Transmitter Interrupt Enable Mode ('TxIntEn')

Value	Action
'1'	Enable Transmitter Interrupt Mode
'0'	Disable Transmitter Interrupt Mode

Table 15. 22 Transmitter Interrupt Enable Mode ('TxIntEn') Field Encoding.

External Status Interrupt Enable ('ExtStatusIntEn')

Value	Action
'1'	Enable External Status Interrupts/Exceptions
'0'	Disable External Status Interrupts/Exceptions

Table 15. 23 External Status Interrupt Enable ('ExtStatusIntEn') Field Encoding.

For additional programming information on the Serial LSB Interrupt Enable Register, refer to Figure 15.13. The bit assignments and field encodings are listed in Table 15. 19, Table 15. 20, Table 15. 21, Table 15. 22, and Table 15. 23.

**Serial Interrupt Vector Register ('SerialIntVectorReg(1:0)')
(Pointer 0x02 for Channel 1 Only)**

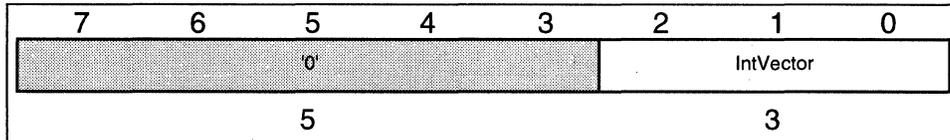


Figure 15.14 Serial Interrupt Vector ('SerialIntVector(1)').

Additional programming information for the Serial Interrupt Vector Register is located in Figure 15.14, Table 15. 24 and Table 15. 25.

Bit	Assignment
7:3	Reserved Low ('0').
2:0	Interrupt Vector

**Table 15. 24 Serial Interrupt Vector Register ('SerialIntVectorReg')
Bit Assignments.**

Value	Action
'7'	Exception Channel 0
'6'	Receiver FIFO Not Empty Channel 0
'5'	External Status Channel 0
'4'	Transmitter FIFO Empty Channel 0
'3'	Exception Channel 0
'2'	Receiver FIFO Not Empty Channel 0
'1'	External Status Channel 0
'0'	Transmitter FIFO Empty Channel 0

Table 15. 25 Interrupt Vector Bit Encoding.

**Serial Receiver Control Register ('SerialRxControlReg(1:0)')
(Write Pointer 0x03)**

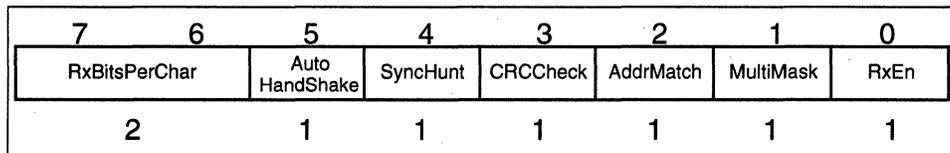


Figure 15.15 Serial Receiver Control Register ('SerialRxControlReg').

Figure 15.15 gives the Serial Receiver Control Registers. Additional programming information is located in Table 15. 26, Table 15. 27, Table 15. 28, Table 15. 29, Table 15. 30, Table 15. 31, Table 15. 32, and Table 15. 33 on page 19.

Bit	Assignment
7:6	Receiver Bits per Character ('RxBitsPerChar')
5	Auto-Handshaking ('AutoHandShake')
4	Sync Char Hunt Mode ('SyncHunt')
3	CRC Checking ('CRCCheck')
2	Address Match Search ('AddrMatch')
1	Address Multiple Match Mask ('MultiMask')
0	Receiver Enable ('RxEn')

Table 15. 26 Serial Receiver Control Register ('SerialRxControlReg') Bit Assignments.

Receiver Bits per Character ('RxBitsPerChar')

Value	Action
'3'	8-bits per character
'2'	6-bits per character
'1'	7-bits per character
'0'	5-bits per character

Table 15. 27 Receiver Bits per Character ('RxBitsPerChar') Field Encoding.

Auto-Handshaking ('AutoHandShake')

This field controls whether handshaking is performed for the Modem signals (CTS, RTS, DCD).

Value	Action
'1'	Enable Auto-Handshaking for Modem Signals
'0'	Disable Auto-Handshaking for Modem Signals

Table 15. 28 Auto-Handshaking ('AutoHandShake') Field Encoding.

Sync Char Hunt Mode ('SyncHunt')

Value	Action
'1'	Enable Sync Char Hunt Mode
'0'	Disable Sync Char Hunt Mode

Table 15. 29 Sync Char Hunt ('SyncHunt') Mode Field Encoding.

CRC Checking ('CRCCheck')

Also see the Serial Transmitter Control Register CRC-16 field.

Value	Action
'1'	Enable CRC Checking on the Receiver
'0'	Disable CRC Checking on the Receiver

Table 15. 30 CRC Checking ('CRCCheck') Field Encoding.

Address Match Search ('AddrMatch')

Value	Action
'1'	Enable Address Match Search
'0'	Disable Address Match Search

Table 15. 31 Address Match Search ('AddrMatch') Field Encoding.

Address Multiple Match Mask ('MultiMask')

In non-SDLC synchronous modes:

Value	Action
'1'	Disable Sync Char Match Load (receive sync char as data)
'0'	Enable Sync Char Stripping

In SDLC mode:

Value	Action
'1'	Address Multi Mask
'0'	Compare all Address bits

Table 15. 32 Address Multiple Match Mask ('MultiMask') Field Encoding.

Receiver Enable ('RxEn')

Value3	Action
'1'	Enable Receiver
'0'	Disable Receiver

Table 15. 33 Enable Receiver ('RxEn') Field Encoding.

**Serial LSB Mode Register ('SerialLSBModeReg(1:0)')
(Write Pointer 0x04)**

7	6	5	4	3	2	1	0
ClkSample		SyncCharMode		StopBit		ParityMode	ParityEn
2		2		2		1	1

Figure 15.16 Serial LSB Mode Register ('SerialLSBModeReg').

The Serial LSB Mode Register fields are given in Figure 15.16. Additional programming information for this register is located in Table 15. 34, Table 15. 35, Table 15. 36, Table 15. 37, Table 15. 38, and Table 15. 39.

Bit	Assignment
7:6	Clock Sampling Mode ('ClkSample')
5:4	Sync Character Mode ('SyncCharMode')
3:2	Stop Bit Mode ('StopBit')
1	Parity Mode ('ParityMode')
0	Parity Enable ('ParityEn') Field Encoding

Table 15. 34 Serial LSB Mode Register ('SerialLSBModeReg') Bit Assignments.

Clock Sampling Mode ('ClkSample')

Value	Action
'3'	%64 Clock Mode
'2'	%32 Clock Mode
'1'	%16 Clock Mode (typical)
'0'	%1 Clock Mode

Table 15. 35 Clock Sampling Mode ('ClkSample') Field Encoding.

Sync Character Mode ('SyncCharMode')

Value	Action
'3'	External Sync
'2'	SDLC Flag
'1'	16-bit Sync
'0'	8-bit Sync

Table 15. 36 Sync Character Mode ('SyncChar') Field Encoding.

Stop Bit Mode ('StopBit')

Value	Action
'3'	2 Stop Bits
'2'	1.5 Stop Bits
'1'	1 Stop Bit
'0'	Sync Mode

Table 15. 37 Stop Bit Mode ('StopBit') Field Encoding.

Parity Mode ('ParityMode')

Value	Action
'1'	Even Parity
'0'	Odd Parity

Table 15. 38 Parity Mode ('ParityMode') Field Encoding.

Parity Enable ('ParityEn')

Value	Action
'1'	Enable Parity
'0'	Disable Parity

Table 15. 39 Parity Enable ('ParityEn') Field Encoding.**Serial Transmitter Control Register ('SerialTxControlReg(1:0)')**
(Write Pointer 0x05)

7	6	5	4	3	2	1	0
DTRN	TxBitsPerChar	TxBreak	TxEn	CRCType	RTSN	TxCRCEn	
1	2	1	1	1	1	1	1

Figure 15.17 Serial Transmitter Control Register ('SerialTxControlReg')

The Serial Transmitter Control Register fields are given in Figure 15.17. Additional programming information for this register is located in Table 15. 40, Table 15. 41, Table 15. 42, Table 15. 43, Table 15. 44, Table 15. 45, Table 15. 46, and Table 15. 47 on page 22.

Bit	Assignment
7	DTR Status ('DTRN')
6:5	Transmitter Bits per Character ('TxBitsPerChar')
4	Transmitter Break ('TxBreak')
3	Transmitter Enable ('TxEn')
2	CRC Type ('CRCType')
1	RTS Status ('RTSN')
0	Transmitter CRC Enable ('TxCRCEn')

Table 15. 40 Serial Transmitter Control Register ('SerialTxControlReg(1:0)') Bit Assignments.**DTR Status ('DTRN')**

Value	Action
'1'	SerialDTR() Signal Low
'0'	SerialDTR() Signal High

Table 15. 41 DTR Status ('DTRN') Field Encoding.

Transmitter Bits per Character ('TxBitsPerChar')

Value	Action
'3'	8-bits per char
'2'	6-bits per char
'1'	7-bits per char
'0'	5-bits per char

Table 15. 42 Transmitter Bits per Character ('TxBitsPerChar').

Transmitter Break ('TxBreak')

Value	Action
'1'	Transmit Break Sequence
'0'	Normal

Table 15. 43 Transmitter Break ('TxBreak') Field Encoding.

Transmitter Enable ('TxEn')

Value	Action
'1'	Enable Transmitter
'0'	Disable Transmitter

Table 15. 44 Transmitter Enable ('TxEn') Field Encoding.

CRC Type ('CRCType')

Value	Action
'1'	Use CRC-16 ($x^{16}+x^{15}+x^2+1$)
'0'	Use CRC-CCITT ($x^{16}+x^{12}+x^5+1$)

Table 15. 45 CRC Type ('CRCType') Field Encoding.

RTS Status ('RTSN')

Value	Action
'1'	Set SerialRTS pin low
'2'	Set SerialRTS pin high

Table 15. 46 RTS Status ('RTSN') Field Encoding.

Transmitter CRC Enable ('TxCRCEn')

Value	Action
'1'	Enable CRC Generator
'0'	Disable CRC Generator

Table 15. 47 Transmitter CRC Enable ('TxCRCEn') Field Encoding.

Serial LSB Sync Character/Address Register
 ('SerialLSBSyncCharReg(1:0)')
 (Pointer 0x06),

Serial MSB Sync Character/Address Register
 ('SerialMSBSyncCharReg(1:0)')
 (Pointer 0x07)

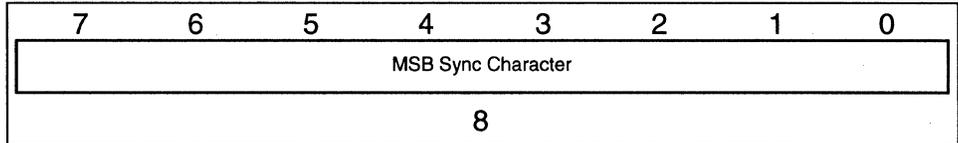


Figure 15.18 Serial MSB Sync Character/Address Register ('SerialMSBSyncCharReg').

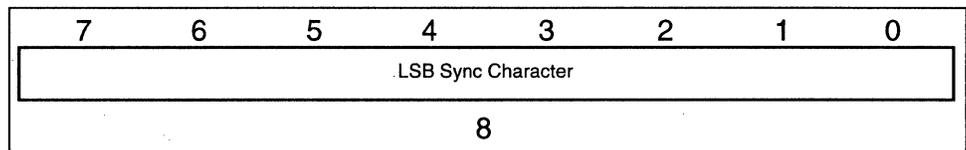


Figure 15.19 Serial LSB Sync Character/Address Register ('SerialLSBSyncCharReg').

Figure 15.18 and Figure 15.19 give the MSB and LSB Sync Character/Address Register fields.

Note: SDLC should always write 0x7E into MSB Sync Register.

6-bit Monosync:

SyncPattern(5:0) -> LSB(5:0)

SyncPattern(1:0) -> LSB(7:6)

SyncPattern(5:0) -> MSB(7:2)

8-bit Monosync, fill both MSB and LSB registers with the same 8-bit value:

SyncPattern(7:0) -> LSB(7:0)

SyncPattern(7:0) -> MSB(7:0)

12-bit Bisync:

'b1111' -> LSB(3:0)

SyncPattern(3:0) -> LSB(7:4)

SyncPattern(11:4) -> MSB(7:0)

16-bit Bisync:

SyncPattern(7:0) -> LSB(7:0)

SyncPattern(15:0) -> MSB(7:0)

SDLC:

Address(7:0) -> LSB(7:0)

'b01111110' -> MSB(7:0)

SDLC Multi:

AddressRange(7:4) -> LSB(7:4)

'b01111110' -> MSB(7:0)

**Transmitter Data FIFO Register ('SerialDataFIFOReg(1:0)')
(Write Pointer 0x08)**

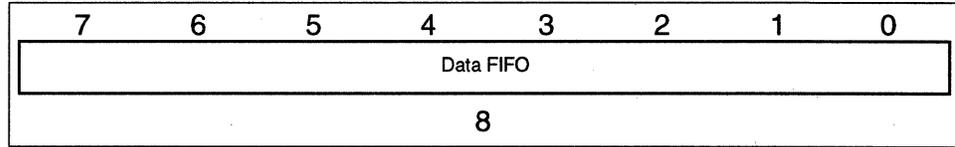


Figure 15.20 Transmitter Data FIFO Register ('SerialDataFIFOReg').

The SerialDataFIFOReg, see Figure 15.20, is also directly accessible by the CPU by writing address 0xFFFF_E800 for Channel 0 or 0xFFFF_E808 for Channel 1.

**Serial Reset Register ('SerialResetReg(1:0)')
(Write Pointer 0x09)**

Programming information for this register is located below in Table 15. 48, Table 15. 49, and Table 15. 50.

Bit	Assignment
7:6	Master Reset ('Reset')
5	Reserved High ('1')
4	Reserved Low ('0')
3	Interrupt Enable ('IntEn')
2:0	Reserved High ('1')

Table 15. 48 Serial Reset Register ('SerialResetReg') Bit Assignments.

Value	Action
'3'	Reset Both Channels (Master Serial Port Reset)
'2'	Reset Channel 0
'1'	Reset Channel 1
'0'	No Effect

Table 15. 49 Master Reset ('Reset') Field Encoding.

Interrupt Enable ('IntEn')

Value	Action
'1'	Global Interrupt Enable Mask
'0'	Global Interrupt Disable Mask

Table 15. 50 Interrupt Enable ('IntEn') Field Encoding.

Serial MSB Mode Register ('SerialModeReg(1:0)') (Write Pointer 0x0A)

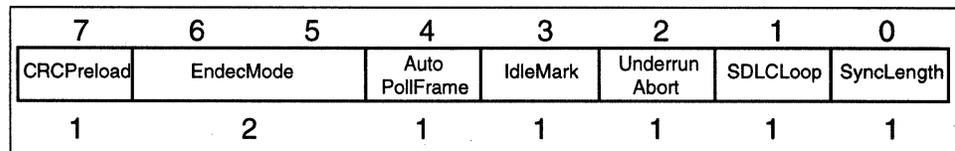


Figure 15.21 Serial MSB Mode Register ('SerialModeReg').

Figure 15.21 shows the Serial MSB Mode Registers. Refer to Table 15. 51, Table 15. 52, Table 15. 53, Table 15. 54, Table 15. 55, Table 15. 56, Table 15. 57, and Table 15. 58 on page 26 for additional programming information.

Bit	Assignment
7	CRC Preload ('CRCPreload')
6:5	Endec Mode ('EndecMode')
4	Automatic SDLC Poll Frame Mode ('AutoPollFrame')
3	Automatic SDLC Transmitter Idle Mark or Flag ('IdleMark')
2	Automatic SDLC Transmitter Underrun Abort ('UnderrunAbort')
1	SDLC Loop Mode ('SDCLoop')
0	Sync Char Length ('SyncLength')

Table 15. 51 Serial MSB Mode Register ('SerialModeReg') Bit Assignments.

CRC Preload ('CRCPreload')

Value	Action
'1'	Preload CRC generator/checker with 1's (Normal).
'0'	Preload CRC generator/checker with 0's.

Table 15. 52 CRC Preload ('CRCPreload') Field Encoding.

Endec Mode ('EndecMode')

Value	Action
'3'	NRZ (typical asynchronous data endec)
'2'	NRZI
'1'	FM1
'0'	FMO

Table 15. 53 Endec Mode ('EndecMode') Field Encoding.

Automatic SDLC Poll Frame Mode ('AutoPollFrame')

Value	Action
1	Send Flag when done
0	Send Flag and revert to 1-bit delay when done

Table 15. 54 Automatic SDLC Poll Frame Mode ('AutoPollFrame') Field Encoding.

Automatic Transmitter SDLC Idle Mark ('IdleMark')

Value	Action
'1'	Send Mark Highs ('1') when transmitter is idle
'0'	Send Flag when transmitter is idle (typical)

Table 15. 55 Automatic Mark or Flag during Transmitter Idle ('IdleMark') Field Encoding.

Automatic Transmitter Underrun Abort ('UnderrunAbort')

Value	Action
'1'	Send Abort and Flag instead of CRC on transmitter underrun
'0'	Send CRC on transmitter underrun

Table 15. 56 Automatic Transmitter Underrun Abort Field Encoding.

SDLC Loop Mode ('SDLCLoop')

Value	Action
'1'	Attach Transmitter Data to Receiver Data
'0'	No Effect

Table 15. 57 SDLC Loop Mode ('SDLCLoop') Field Encoding.

Sync Char Length ('SyncLength')

Value	Action
'1'	6-bit sync character length for Monosync. 12-bit Receiver sync character length and 16-bit transmitter sync character length for Bisync
'0'	8-bit Sync character length

Table 15. 58 Sync Char Length ('SyncLength') Field Encoding.

**Serial LSB Clock Control Register ('SerialLSBClkControlReg(1:0)')
(Write Pointer 0x0B)**

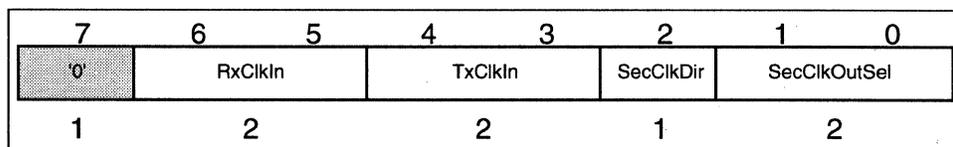


Figure 15.22 Serial LSB Clock Control Register ('SerialLSBClkControlReg').

Figure 15.22 gives the fields for the Serial LSB Clock Control Register. For additional programming information, refer to Table 15. 59, Table 15. 60, Table 15. 61, Table 15. 62, Table 15. 62, and Table 15. 63.

Bit	Assignments
7	Reserved Low ('0')
6:5	Receiver Clock Input Source ('RxClkIn')
4:3	Transmitter Clock Input Source ('TxClkIn')
2	SerialSecondaryClk() Direction ('SecClkDir')
1:0	SerialSecondaryClk() Output Clock Select ('SecClkOutSel')

Table 15. 59 Serial LSB Clock Control Register ('SerialLSBClkControlReg') Bit Assignments.

Receiver Clock Input Source ('RxClkIn')

Value	Action
'3'	Use PLL
'2'	Use Baud Rate Clock Generator
'1'	Use SerialSecondaryClk()
'0'	Use SerialPrimaryClkIn()

Table 15. 60 Receiver Clock Input Source ('RxClkIn') Field Encoding.

Transmitter Clock Input Source ('TxClkIn')

Value	Action
'3'	Use PLL
'2'	Use Baud Rate Clock Generator
'1'	Use SerialSecondaryClk()
'0'	Use SerialPrimaryClkIn()

Table 15. 61 Transmitter Clock Input Source ('TxClkIn') Field Encoding.

SerialSecondaryClk() Direction ('SecClkDir')

Value	Action
'1'	SerialSecondaryClk() is an output
'0'	SerialSecondaryClk() is an input

Table 15. 62 SerialSecondaryClk() Direction ('SecClkDir') Field Encoding.

SerialSecondaryClk() Output Clock Select ('SecClkSel')

If the SerialSecondaryClk() is configured as an output both in the SerialSecondaryClk() Direction ('SecClkDir') Field and in the corresponding PIO Field, then the output can be 1 of 4 choices.

Value	Action
'3'	Use PLL
'2'	Use Baud Rate Clock Generator
'1'	Use SerialSecondaryClk()
'0'	Use SerialPrimaryClkIn()

Table 15. 63 SerialSecondaryClk() Output Clock Select ('SecClkOutSel') Field Encoding (if output is enabled).

Serial LSB Baud Rate Clock Generator Compare Register ('SerialLSBBaudCompareReg(1:0)') (Pointer 0x0C),

Serial MSB Baud Rate Clock Generator Compare Register ('SerialMSBBaudCompareReg(1:0)') (Pointer 0x0D)

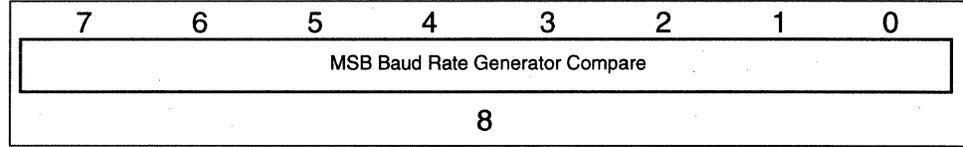


Figure 15.23 Serial MSB Baud Rate Clock Generator Compare Register ('SerialMSBBaudCompareReg').

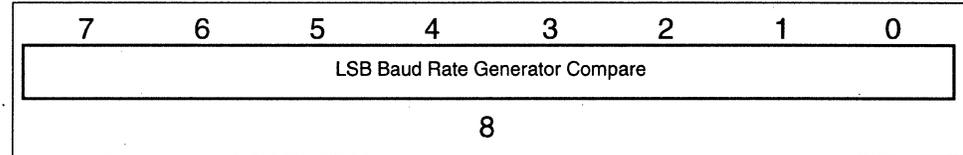


Figure 15.24 Serial LSB Baud Rate Clock Generator Compare Register ('SerialLSBBaudCompareReg').

The Serial LSB and MSB Baud Rate Clock Generator Compare Register fields are shown in Figure 15.23 and Figure 15.24.

If the Baud Rate Clock Generator is used, then the Compare value must be programmed to:

$$\mathbf{BRCGCompareReg = [BRCGClkIn / (2 * BaudRate * ClkSample)] - 2.}$$

Common values for a SysClk Baud Rate Clock Generator Clock Input with Clock Sampling Mode of %16 are:

18.42MHz		
Baud	Compare	% Error
19200	28	0
9600	58	0
2400	238	0
20.0MHz		
Baud	Compare	% Error
9600	63	0.160
25.0MHz		
Baud	Compare	% Error
9600	79	0.469
33.0MHz		
Baud	Compare	% Error
9600	105	0.394

PLL Control Register ('PLLControlReg(1:0)')
(Write Pointer 0x0E)

7	6	5	4	3	2	1	0
PLLMode			Loopback	Echo	'1'	BRCGCkIn	BRCGE _n
3			1	1	1	1	1

Figure 15.25 PLL Control Register ('PLLControlReg').

Figure 15.25 gives the fields for the PLL Control Register. Additional programming information for this register is located in Table 15. 64, Table 15. 65, Table 15. 66, Table 15. 67, Table 15. 68, and Table 15. 69.

Bit	Assignment
7:5	PLL Mode ('PLLMode')
4	Loopback Test Mode ('Loopback')
3	Echo Test Mode ('Echo')
2	Reserved High ('1')
1	Baud Rate Clock Generator Clock Input ('BRCGCkIn')
0	Baud Rate Clock Generator Enable ('BRCGE _n ')

Table 15. 64 PLL Control Register ('PLLControlReg').

PLL Mode ('PLLMode')

Value	Action
'7'	Use NRZI for PLL ClkIn
'6'	Use FM for PLL ClkIn
'5'	Use $\overline{\text{SerialPrimaryClkIn}}[0]$ for PLL ClkIn
'4'	Use Baud Rate Clock Generator PLL ClkIn
'3'	Disable PLL
'2'	Reset Missing Clock Status
'1'	PLL Search Mode
'0'	No action

Table 15. 65 PLL Mode ('PLLMode') Field Encoding.

Loopback Test Enable ('Loopback')

The loopback test provides a manufacturing test mode where the transmitter data path is tied directly to the receiver data path such that the external SerialTxData and SerialRxData pins are bypassed.

Value	Action
'1'	Enable Loopback Test
'0'	Normal Mode

Table 15. 66 Loopback Test Enable ('Loopback') Field Encoding.

Echo Test Enable ('Echo')

Value	Action
'1'	Enable Echo Test Mode
'0'	Normal Mode

Table 15. 67 Echo Test Enable ('Echo') Field Encoding.

Baud Rate Clock Generator Clock Input ('BRGCkIn')

Value	Action
'1'	Use $\overline{\text{SysClk}}$ as the input clock
'0'	Use $\overline{\text{SerialPrimaryClkIn}}$ as the input clock

Table 15. 68 Baud Rate Clock Generator Clock Input ('BRGCkIn') Field Encoding.

Baud Rate Clock Generator Enable ('BRCGEn')

Value	Action
'1'	Enable the Baud Rate Clock Generator counter
'0'	Disable the Baud Rate Clock Generator counter (default)

Table 15. 69 Baud Rate Clock Generator Enable ('BRCGEn') Field Encoding.

**Serial MSB Interrupt Enable Register ('SerialIntEnReg(1:0)')
(Write Pointer 0x0F)****Serial MetaMetaRegister ('SerialMetaMetaReg(1:0)')
(Write Pointer 0x0F)**

7	6	5	4	3	2	1	0
BreakIntEn	Underrun IntEn	CTSIntN	SyncIntEn	DCDNIntEn	FrameFIFO En	ZerolntEn	MetaMeta RegEn
1	1	1	1	1	1	1	1

Figure 15.26 Serial MSB Interrupt Enable Register ('SerialIntEnReg').

Figure 15.26 gives the fields for the Serial MSB Interrupt Enable Register. Additional programming information for this register is located in Table 15. 70, Table 15. 71, and Table 15. 72.

Bit	Assignment
'7'	Break or Abort Interrupt Enable ('BreakIntEn')
'6'	Transmitter Underrun or EOM Condition Interrupt Enable ('UnderrunIntEn')
'5'	CTS Interrupt Enable ('CTSIntEn')
'4'	Sync Char Interrupt Enable ('SyncIntEn')
'3'	DCD Interrupt Enable ('DCDIntEn')
'2'	Frame Status FIFO Enable (MetaMetaRegister B) Enable ('FrameFIFOEn')
'1'	Baud Rate Clock Generator Zero Condition Interrupt Enable ('ZeroIntEn')
'0'	MetaMeta Control Register Enable ('MetaMetaRegEn')

Table 15. 70 Serial MSB Interrupt Enable Register ('SerialIntEnReg') Bit Assignments.

Interrupt Enable Bits

Note that interrupts feed into a common internal interrupt line, SerialIntN, that is further controlled by the R36100's Interrupt Controller. Thus to be fully enabled, the corresponding Interrupt Enable bit must be set, the master Interrupt Enable bit in the Serial Reset Register must be set, the Interrupt Controller SerialIntN IE bit must be set, the CPU Core Status Register IE bit for SysExcInt5 must be set, and finally, the CPU Core Status Register master IE bit must be set.

Value	Action
'1'	Enable Interrupt
'0'	Disable Interrupt

Table 15. 71 Interrupt Enable Bit Field Encoding.

MetaMeta Control Register Enable ('MetaMetaRegEn')

Value	Action
'1'	Enable writes to the Meta Meta Control Register (Write Pointer 0x07).
	Enable reads from the Meta Meta MSB and LSB Frame FIFO Status Registers (Read Pointers 0x07 and 0x06)
'0'	Enable writes to the MSB Sync Char Register (Pointer 0x07)

Table 15. 72 MetaMeta Control Register Enable ('MetaMetaRegEn').

**Serial LSB Frame Status FIFO Data Register
('SerialLSBFrameStatusDataReg(1:0)') (MetaPointer 0x06),**

**Serial MSB Frame Status FIFO Data Register
('SerialMSBFrameStatusDataReg(1:0)') (MetaPointer 0x07)**

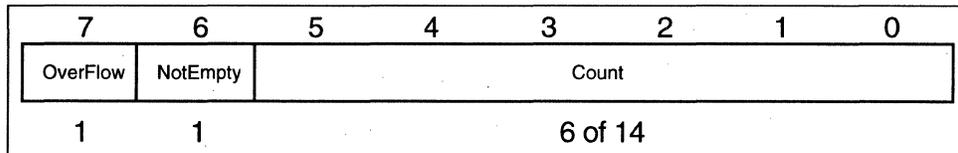


Figure 15.27 Serial MSB Frame Status FIFO Data Register ('SerialMSBFrameStatusDataReg').

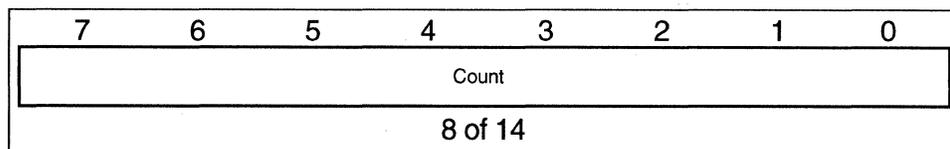


Figure 15.28 Serial LSB Frame Status FIFO Data Register ('SerialLSBFrameStatusDataReg').

Figure 15.27 and Figure 15.28 show the fields for the Serial MSB and LSB Frame Status FIFO Data Registers. Refer to Table 15. 73, Table 15. 74, and Table 15. 75 for additional programming instructions.

Bit	Assignments
MSB 7	Frame Status FIFO Overflow Flag ('Overflow')
MSB 6	Frame Status Not Empty Flag ('NotEmpty')
MSB 5:0 LSB 7:0	Count

Table 15. 73 Serial LSB and MSB Frame Status FIFO Data Register ('SerialMSB-FrameStatusDataReg' and ('SerialMSBFrameStatusDataReg') Bit Assignments.

Frame Status FIFO Overflow Flag ('Overflow')

Value	Action
'1'	Frame Status FIFO Overflow Flag
'0'	Frame Status FIFO Okay Flag

Table 15. 74 Frame Status FIFO Overflow Flag ('Overflow').

Frame Status Not Empty Flag ('NotEmpty')

Value	Action
'1'	Frame Status FIFO Not Empty (Entry is Ready) Flag
'0'	Frame Status FIFO Empty Flag

Table 15. 75 Frame Status Not Empty Flag ('NotEmpty').

Count ('Count')

The Count Field holds the number of byte data received during a frame.

**Serial Meta Meta Register ('SerialMetaMetaReg(1:0'))
(Write Pointer 0x07)**

7	6	5	4	3	2	1	0
0	1	WaitCRC	0	Auto TxDataHigh	Auto SDLCRTS	Auto TxUnderrun	Auto TxFlag
1	1	1	1	1	1	1	1

Figure 15.29 Serial Meta Meta Register ('SerialMetaMetaReg').

Figure 15.29 shows the fields of the Serial Meta Meta Register. Additional programming instructions for this register are located in Table 15. 76, Table 15. 77, Table 15. 78, Table 15. 79, Table 15. 80, and Table 15. 81.

Bit	Assignments
7	Reserved Low ('0')
6	Reserved High ('1')
5	Wait for Receiver CRC ('WaitCRC')
4	Reserved Low ('0')
3	Automatically Drive SerialTxData High when disabled ('AutoTxDataHigh')
2	Automatically De-assert SerialRTS after SDLC Flag ('AutoSDLCRTS')
1	Automatically Reset Transmitter Underrun ('AutoTxUnderrun')
0	Automatic Insertion of Transmitter Opening Flag ('AutoTxFlag')

Table 15. 76 Serial Meta Meta Register ('SerialMetaMetaReg') Bit Assignments.

Wait for Receiver CRC ('WaitCRC')

Value	Action
'1'	Wait for CRC to completely shift into the receiver before loading into the Receiver Data FIFO.
'0'	Don't wait for CRC (CRC will be incorrect).

Table 15. 77 Wait for Receiver CRC ('WaitCRC') Field Encoding.

Automatically Drive SDLC SerialTxData High when disabled ('AutoTxDataHigh')

Value	Action
'1'	In the SDLC NRZI mode, automatically drive SerialTxData High when the transmitter is disabled.
'0'	No special effect.

Table 15. 78 Automatically Drive SerialTxData High when disabled ('AutoTxDataHigh') Field Encoding.

Automatically De-assert SerialRTS after SDLC Flag ('AutoSDLCRTS')

Value	Action
'1'	In the SDLC mode, automatically de-assert SerialRTS high after the SDLC Flag at the end of the frame completes.
'0'	No special effect.

Table 15. 79 Automatically De-assert SerialRTS after SDLC Flag ('AutoSDLCRTS') Field Encoding.

Automatically Reset Transmitter Underrun ('AutoTxUnderrun')

Value	Action
'1'	In Synchronous modes such as SDLC, automatically reset the transmitter underrun latch when the first byte of the frame is sent.
'0'	No special effect.

Table 15. 80 Automatically Reset Transmitter Underrun ('AutoTxUnderrun') Field Encoding.

Automatic Insertion of Transmitter Opening Flag ('AutoTxFlag')

Value	Action
'1'	In Synchronous modes such as SDLC, automatically insert an opening flag before transmitting the first data character of a frame.
'0'	No special effect.

Table 15. 81 Automatic Insertion of Transmitter Opening Flag ('AutoTxFlag') Field Encoding.



Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. One of the many on-chip peripherals is the *Bidirectional Centronics Parallel Port*, as described in this chapter.

This chapter will provide an overview of the Bidirectional Centronics Parallel Port register interface, a description of the signal pins, and various aspects of the signal timing.

Features

- Bidirectional ParallelPortTarget/Peripheral/PrinterController provided on-chip
- Provides 9 pin interface to Bidirectional Centronics IEEE 1284 Standard Parallel Port
- Provides 2 pins for host transceiver control
- Reuses 3 I/O Controller pins for peripheral transceiver control
- Uses external transceiver or bi-directional FIFO for data storage
- DMA auto-initiate via internal interrupt
- Compatible 8-bit input host to peripheral protocol (backward compatibility with Centronics standard)
- Nibble mode peripheral to host output protocol (Microsoft/PC standard)
- Byte mode peripheral to host output protocol (IBM PS2 applications)
- ECP bidirectional protocol (emerging Windows PC/Laser standard)
- EPP bidirectional protocol (PC/Datacom applications)
- 200KBytes/sec to 1 MByte/sec

Block Diagram

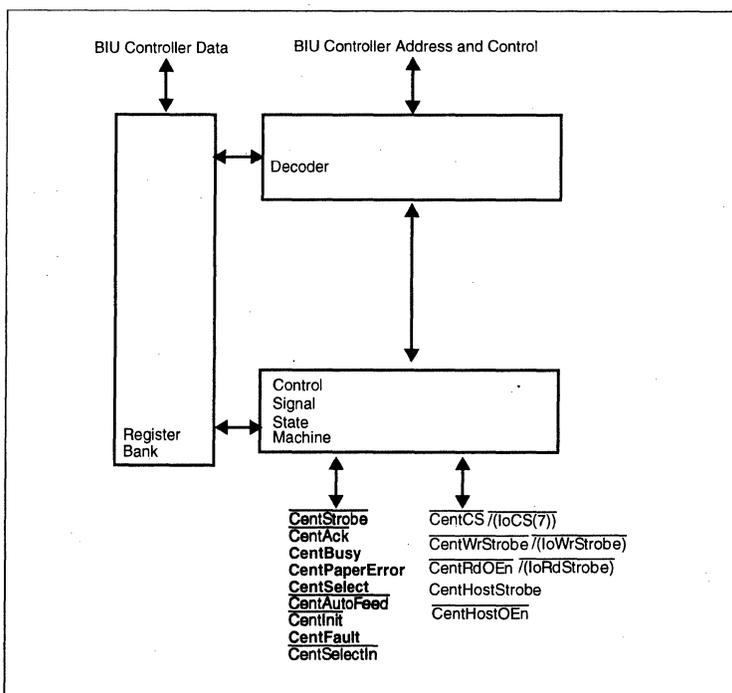


Figure 16.1 Block Diagram of the Bidirectional Parallel Port.

Overview

The Bidirectional Parallel Port Target/Peripheral/Printer Interface is an implementation of the interface described in "Standard Signaling Method for a Bidirectional Parallel Peripheral Interface for Personal Computers, IEEE Standard 1284.

The Bidirectional Parallel Port Interface's overall effect is to allow Laser Printers or add-in communication cards (such as external SCSI drives or external Ethernet ports) to communicate with a PC host in both directions by use of receive and transmit channels.

When using the R36100 for an Apple laser printer peripheral, information commonly associated with Apple LocalTalk/AppleTalk functionality, such as printer status and resident font information, is passed back from the printer to the host Macintosh. However, in the IBM PC compatible environment, the original Centronics compatible mode is only unidirectional and cannot report general purpose status information back to the PC host.

With the addition of a reverse transfer 4-bit IEEE 1284 nibble mode, the Centronics port on the printer peripheral can now communicate bidirectionally with the majority of legacy PC hosts by using the present printer status lines to pass 4-bits at a time back to the PC host. In newer PC's, such as the PS2 series, the PC hosts can use a truer bidirectional mode such as the IEEE 1284 reverse transfer byte mode.

The R36100 also supports the newer IEEE 1284 Extended Capabilities Port mode (ECP) and IEEE 1284 Enhanced Parallel Port (EPP) mode, which provide more efficient interlocked handshaking as well as symmetric byte and host controlled read/write byte channel protocols, respectively. Both ECP and EPP are commonly found on Enhanced IDE I/O PC cards.

The Bidirectional Parallel Port Interface uses 14 pins (see block diagram in Figure 16.1). The pins include 9 control signals multiplexed in/out with PIO. The data lines are supported by an external 8-bit data register transceiver chip or bidirectional FIFO controlled by any one of the I/O Controller chip selects, the I/O read strobe, and the I/O write strobe. The pins also include 2 external register control lines. One of them clocks the data from the PC to the printer port. The other line output enables the external register to the PC.

When used with two 8-bit external buffers/transceivers and a compliant physical connector, the R36100 Bidirectional Parallel Port Interface implementation meets the IEEE 1284 definition of a compliant device. The R36100 supports the following peripheral modes:

- Compatible (standard forward transfer)
- nibble (4-bit reverse transfer)
- byte (8-bit reverse transfer)
- ECP (Extended Capabilities Port) (forward and reverse interlocked handshake transfers with arbitration for host/port control)
- EPP (Enhanced Parallel Port) (host controlled forward and reverse read/write-like transfers)

The R36100 also contains support for the negotiation phase necessary for transition between the different modes. As described below, each mode has different phases associated with them:

Compatible Mode Phases are:

1. Forward Data Transfer
2. Forward Idle
3. (Negotiation)

Nibble Mode and Byte Mode Phases are:

1. Forward Data Transfer
2. Forward Idle
3. Negotiation
4. Host Busy Data Available
5. Reverse Data Transfer
6. Host Busy Data Not Available
7. Reverse Idle
8. Interrupt Host
9. Terminate

Note: In nibble and byte modes, the R36100 Centronics port always goes from state 3 --> 7 and then to state 8 --> 4, never from state 3 to 4. For data ready status, this requires the host to poll using the 7/4 states, not in the 3/4 states.

ECP Mode Phases are:

1. Forward Data Transfer
2. Forward Idle
3. Negotiation
4. Setup
5. Forward Idle
6. Forward
7. Forward to Reverse
8. Reverse Idle
9. Reverse
10. Reverse to Forward
11. Terminate

EPP Mode Phases are:

1. Forward Data Transfer
2. Forward Idle
3. Negotiation
4. Initial EPP Idle
5. Address Read
6. Data Read
7. Address Write
8. Data Write
9. EPP Idle
10. Terminate

Support for the compatible mode includes the three variations listed in Table 16.1.

Variation	Strobe/Busy	Busy/Ack
Centronics Classic	Busy-after-Strobe	Ack(2500 ns)-after-Busy
IBM/Epson	Busy-after-Strobe	Ack(2500 ns)-while-Busy
Standard 1284	Busy-while-Strobe	Ack(500 ns)-in-Busy

Table 16.1 Compatible Forward Data Transfer Variations.

Negotiation Phase

The Parallel Port Interface is initially put into "compatible mode" after reset. While in compatible mode, the host can send data to the peripheral in a forward data transfer phase. In order to get into any of the other modes which support reverse data transfers, the port must undergo a negotiation phase in order to see if the port can support the requested mode. The Bidirectional Parallel Port Interface software driver must also configure the compatible mode to one of the three supported modes (IBM, classic, or standard) and to a data transfer option (DMA or interrupt per byte). Setting the modes and options is done by writing to the mode register.

In the interrupt per byte mode, the R36100 will read data from external Centronics Data Register each time it responds to a $\overline{\text{CentRdInt}}$ interrupt. In the DMA mode, the R36100 will initialize one of the Internal DMA Channel Controllers register to the start of the DMA operation. The Bidirectional Parallel Port Interface software driver can be notified by interrupt when the DMA counter reaches the terminal count.

The negotiation is indicated by:

1. Host asserts 1284Active (nSelectIn) and de-asserts HostBusy (nAutoFd).
2. The peripheral responds by bringing AckDataReq (PErr), nDataAvail (nFault), Xflag (Select) high and PtrClk (nAck) low.
3. Host nStrobes 8-bit extensibility request value on the data lines and also brings HostBusy (nAutoFd) high.
4. Peripheral sets Xflag (Select) to a particular value, and in the nibble and byte modes, nDataAvail (nFault) and AckDataReq (PErr). Busy and Ptr (nAck) are set appropriately.

After step 1, the R36100 is interrupted by the Parallel Port Interface, $\overline{\text{CentWrInt}}$. The interrupt service routine must then read the extensibility request value from the external Centronics Data Register and write the appropriate mode and response bits back to the Parallel Port Interface so that it can finish the negotiation. If the extensibility link bit is asserted, then a second $\overline{\text{CentWrInt}}$ will occur during the negotiation.

A host request to return to compatibility mode, from any of the other modes, is indicated to the R36100 by the assertion of the $\overline{\text{CentRst}}$ interrupt.

Nibble Mode Phase

The Parallel Port Interface will interrupt the R36100 by asserting $\overline{\text{CentWrInt}}$ when the host requests a two-nibble (8-bits total) transfer. The R36100 will respond by writing data to the Parallel Port's Nibble Data Register. The Parallel Port Interface sends the two nibbles to the host over the appropriate Centronics control lines in two consecutive nibble transactions.

Byte Mode Phase

The Parallel Port Interface will interrupt the R36100 by asserting $\overline{\text{CentWrInt}}$ when the host requests a byte transfer. The R36100 will respond by writing data to the external Centronics Data Register.

Note: In Nibble and Byte Mode, the peripheral can arbitrate for the port, but only if it is left in reverse idle phase.

Extended Capabilities Port (ECP) Mode Phase

The ECP Mode allows both the host and the printer port to arbitrate for the bus and send commands/data to each other. Up to 128 different channels (communication streams) are supported by the protocol.

DMA and interrupt-per-byte options are supported for the ECP mode.

In the interrupt-per-byte option, the Parallel Port Interface will first assert $\overline{\text{CentRdInt}}$ for host read or write requests, and then it will assert $\overline{\text{CentWrInt}}$ for host write requests or $\overline{\text{CentRdInt}}$ for host read requests. The R36100 will read or write from the external Centronics Data Register in response to the interrupts.

In reverse transfer, in response to $\overline{\text{CentWrInt}}$, the R36100 needs to write to the Parallel Port's Status Register (to the Busy bit) to indicate if it is sending a command or data byte, and also write the data/command to an external Centronics Data Register.

In forward transfer, in response to $\overline{\text{CentRdInt}}$, the R36100 needs to read from the Parallel Port's Command Register (nAutoFd bit) to see if the Centronics Data Register has a data or command byte. Run Length Encoding (RLE) Compression/decompression, if implemented, must be done by the software driver.

In the DMA transfer option, data will be transferred by an internal DMA channel as long as the direction of the host requests matches the direction of the DMA. Software must handle Centronics interrupts until the address and control is set up. Afterwards, a data stream can be handled by DMA. $\overline{\text{CentWrInt}}$ will be asserted when the host requests data. $\overline{\text{CentRdInt}}$ will be asserted when the host sends data or when the host sends a command byte.

Enhanced Parallel Port (EPP) Mode Phases

The EPP mode allows the host to address the printer port much like a read and writable memory interface. However, as per the IEEE1284 specification, the peripheral can not initiate transfers in this mode.

DMA and interrupt-per-byte options are supported for the EPP mode. In the interrupt-per-byte option, the Parallel Port Interface will assert $\overline{\text{CentRdInt}}$ for host read requests, and will assert $\overline{\text{CentWrInt}}$ for host write requests. The R36100 will read or write from the external Centronics Data Register in response to the interrupts.

$\overline{\text{CentWrInt}}$ will be asserted when the host requests data (from the IEEE1284 port to the host). $\overline{\text{CentRdInt}}$ will be asserted when the host sends data or when the host sends an address byte. Software must handle Centronics interrupts until the address and control is set up. Afterwards, a data stream can be handled by DMA.

CPU Control Reserved Mode Phases

This mode enables the CPU to send the lower 5 bit values of the Parallel Port Interface's Status Register.

Programmable Timing

To allow for higher data rates than those specified by the IEEE1284 Standard, the minimum delay on Strobe/Busy and Busy/Ack can be programmed to lower values than the minimum required by the standard. Table 16.2 lists the three interrupt signals and descriptions.

Interrupt Signal Name	Description
CentWrInt	1 clock pulse assertion after host write and for other forward transfer phases.
CentRdInt	Pulse assertion after read or reverse transfer phases.
CentResetInt	1 clock pulse assertion after termination phase and after negotiation phase is complete.
CentNegInt	Asserts during negotiation phase

Table 16.2 Interrupt Descriptions

Pin Descriptions

Note that the following pin descriptions are given in terms of the Centronics Modes. Actually, each mode has various phases that may further define the functionality of the signal. Please refer to IEEE 1284 Standard for additional detail.

Bidirectional Parallel Port Centronics Interface Signals

CentStrobe **Input**
(Aliases: nStrobe, nStrobe, HostClk, nWrite)
 Centronics Strobe:
Compatibility: Data strobe.
Nibble: Acknowledges reverse data transfer.
Byte: Acknowledges reverse data transfer.
ECP: Handshakes with Busy.
EPP: Indicates Address write or Data write.

CentAck **Output**
(Aliases: nAck PtrClk, nAck PtrClk, PeriphClk, Int)
 Centronics Acknowledge:
Compatibility: Data Acknowledge.
Nibble: Data Acknowledge.
Byte: Data Acknowledge.
ECP: Handshakes with HostAck.
EPP: Active High Interrupt.

CentBusy **Output**
(Aliases: Busy PtrBusy, Busy PtrBusy, PeriphAck, nWait)
 Centronics Busy:
Compatibility: Active high indication that the peripheral is busy.
Nibble: In later phases of nibble mode, Data bit 3 and 7.
Byte: Active high indication that the peripheral is busy.
ECP: Flow control in the forward direction, Command/Data bit in the reverse direction.
EPP: Active low wait signal delaying an address or data.

CentPaperError **Output**
(Aliases: PError, PError, AckDataReq, nAckReverse)
 Centronics PaperError:
Compatibility: When asserted with nFault, indicates a Paper Error. Additional uses during other phases including 1284Support.
Nibble: Data bits 2 and 6.
Byte: Same as nFault.
ECP: Request nReverseRequest.
EPP: User Defined (unused).

CentSelect **Output**
(Aliases: Select Xflag, Select Xflag, Xflag)

Centronics Select:

Compatibility: Peripheral is on line. In other phases, the Extensibility Flag (XFlag).

Nibble: Data bits 1 and 5.

Byte: Peripheral is on line. In other phases, the Extensibility Flag (XFlag).

ECP: In some phases, the Extensibility Flag (XFlag).

EPP: User Defined (unused).

CentAutoFeed **Input**
(Aliases: nAutoFd HostBusy, nAutoFd HostBusy, HostAck, nDStrb)

Centronics Auto Feed:

Compatibility: Typically indicates auto linefeed mode, but often unused or redefined. Also used during Negotiation Phase as HostBusy.

Nibble: Typically indicates auto linefeed mode. In other phases, used for several purposes.

Byte: Typically indicates auto linefeed mode. In other phases, used for several purposes.

ECP: Handshakes with PeriphClk.

EPP: Denotes data cycle.

CentInit **Input**
(Aliases: nInit, nInit, nReverseRequest, nInit)

Centronics Initialize:

Compatibility: When pulsed with 1284Active de-asserted, resets to idle phase.

Nibble: When pulsed with 1284Active de-asserted, resets to idle phase.

Byte: When pulsed with 1284Active de-asserted, resets to idle phase.

ECP: Host allows the peripheral to drive the bi-directional data signals.

EPP: When asserted, resets to compatibility mode.

CentFault **Output**
(Aliases: nFault nDataAvail, nFault nDataAvail, nPeriphRequest)

Centronics Fault:

Compatibility: Set low indicating an error. In other phases, set high to ack 1284, data available.

Nibble: Data bits 0 and 4.

Byte: Set low indicating an error. In other phases, set high to ack 1284 and data available.

ECP: Peripheral requests communication with host which host may chose to ignore.

EPP: User defined (unused).

CentSelectIn **Input**
(Aliases: nSelectIn 1284Active, nSelectIn 1284Active, 1284Active, nAStrb)

Centronics Select Input:

Compatibility: Selects this peripheral (if Centronics is shared). In some phases, indicates 1284Active.

Nibble: Selects this peripheral (if Centronics is shared). In some phases, indicates 1284Active.

Byte: Selects this peripheral (if Centronics is shared). In some phases, indicates 1284Active.

ECP: Active high, if de-asserted, return to compatibility mode.

EPP: Indicates an address cycle.

Bidirectional Parallel Port Centronics Peripheral and Host Interface Signals

CentCS() Output

Centronics I/O Chip Select: Use 1 of the 8 $\overline{\text{IoCS}}()$ pins to create this signal. This active low signal is used by the R36100 to select the externally provided 8-bit Centronics Data Register/Transceiver. With some types of transceivers $\overline{\text{CentCS}}()$ must be externally gated with $\overline{\text{CentRdStrobe}}$ and or $\overline{\text{CentWrStrobe}}$.

CentWrStrobe Output

Centronics Write Strobe: Use $\overline{\text{IoWrStrobe}}$ to create this signal. This active low signal is used by the R36100 to write data into a registered transceiver so that the host may later retrieve the data. The transceiver must also be gated with an appropriately programmed $\overline{\text{IoCS}}()$.

CentRdOEn Output

Centronics Read Output Enable: Uses $\overline{\text{IoRdStrobe}}$ to create this signal. This active low signal is used by the R36100 to read data into the peripheral from the registered transceiver which the host had previously stored. The transceiver must also be gated with an appropriately programmed $\overline{\text{IoCS}}()$.

CentHostStrobe Output

Centronics Host Strobe: Similar to $\overline{\text{CentStrobe}}$ but active high and gated for actual host data writes since $\overline{\text{CentStrobe}}$ is also used by various IEEE 1284 modes to acknowledge actions other than writes. Active high output is attached to an external registered transceiver in order to clock/latch-enable the data from the host into the registered transceiver. The CentHostStrobe pin is multiplexed with a PIO pin, and thus the PIO pin must be programmed to the CentHostStrobe special effect and to be an output.

CentHostOEn Output

Centronics Host Output Enable: Active low output attached to an external registered data transceiver in order to allow the host to read data from the registered transceiver. The $\overline{\text{CentHostOEn}}$ pin is multiplexed with a PIO pin, and thus the PIO pin must be programmed to use the $\overline{\text{CentHostOEn}}$ special effect and to be an output.

Register Definitions

Table 16.3 lists the Bidirectional Parallel Port Interface addresses and descriptions. Note that Big Endian software must offset these addresses by b'10 (0x2).

Phys. Address	Description
0xFFFF_EC00	Centronics Sub Mode Register
0xFFFF_EC04	Centronics Status Register
0xFFFF_EC08	Centronics Control Register
0xFFFF_EC0C	Centronics Nibble Data Register
0xFFFF_EC10	Centronics Host Status Register
0xFFFF_EC14	Centronics Minimum Delay Register

Table 16.3 Bidirectional Parallel Port Interface Centronics Controller Registers.

Centronics Sub Mode Register ('CentSubModeReg')

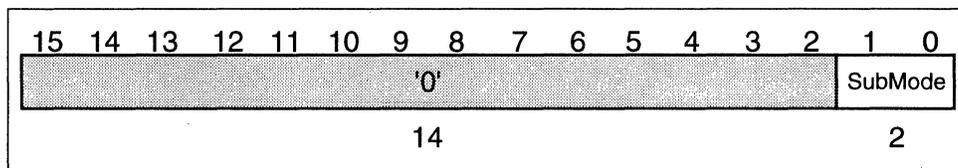


Figure 16.2 Centronics Sub ModeRegister ('CentSubModeReg').

Additional programming information and instructions for the Centronics Sub Mode Register are located in Figure 16.2, Table 16.4, and Table 16.5.

Bit	Description
1:0	CMode

Table 16.4 Centronics Sub ModeRegister ('CentSubModeReg') Bit Assignments.

Centronics Compatible Sub Modes ('SubMode') Field

Value	Action
'11'	Classic Centronics.
'10'	Reserved.
'01'	IBM/Epson.
'00'	Standard 1284 (default).

Note: See Table 16.1 for more information.

Table 16.5 Centronics Compatible Sub Mode ('SubMode') Field Encoding.

Centronics Status Register ('CentStatusReg')

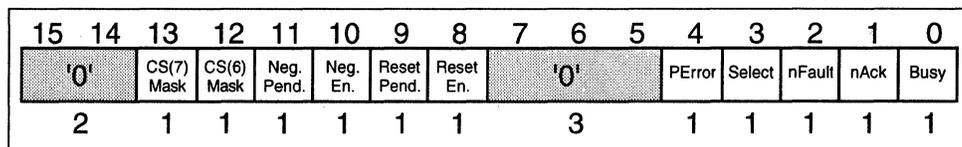


Figure 16.3 Centronics Status Register ('CentStatusReg').

The Centronics Status Register is shown in Figure 16.3. Additional programming information and instructions for this register are located in Table 16.6, Table 16.7, Table 16.8, Table 16.9, Table 16.10, Table 16.11, Table 16.12, Table 16.13, Table 16.14, Table 16.15, and Table 16.16.

Bit	Description
13	IoCS(7) Mask Enable
12	IoCS(6) Mask Enable
11	Negotiation Interrupt Pending
10	Negotiation Interrupt Enable
9	Reset Interrupt Pending
8	Reset Interrupt Enable
4	Printer Error (PErrror)
3	Printer On Line Select (Select)
2	Printer Fault Negated (nFault)
1	Printer is Acknowledging Negated (nAck)
0	Printer is Busy (Busy)

Table 16.6 Centronics Status Register ('CentStatusReg') Bit Assignments.

IoCS(7), IoCS(6) Mask Enable

Typically, a system using the R36100 bi-directional centronics interface will use an external buffer to buffer data between the CPU bus and the Centronics port. These bits control whether either or both of IoCS(7:6) are used to control that buffer. The most common strategy is to use one IoCS for reads and one for writes, although other systems may just use one IoCS and decode read or write from the control bus.

In addition, address bit 15 must be low when accessing the external Centronics data register.

Value	Action
'1'	Use IoCS for Centronics
'0'	Don't use IoCS for Centronics

Table 16.7 IoCS(7:6) Mask Enable Field Encoding

Centronics Negotiation Interrupt Pending

Value	Action
'1'	On reads, means negotiation interrupt is pending.
	On writes, clears the pending interrupt.
'0'	On reads, means no negotiation interrupt is pending.
	On writes, means do not change current interrupt state.

Table 16.8 Centronics Negotiation Interrupt Pending Field.

Centronics Negotiation Interrupt Enable

Value	Action
'1'	Signal Pending Negotiation Interrupts to the interrupt controller.
'0'	Do not signal the interrupt controller (default).

Table 16.9 Centronics Negotiation Interrupt Enable Field.

Centronics Reset Interrupt Pending

Value	Action
'1'	On reads, means reset interrupt is pending.
	On writes, clears the pending interrupt.
'0'	On reads, means no reset interrupt is pending.
	On writes, means do not change current interrupt state.
Note: Must first be written high and then written again low. This differs from the Expansion Interrupt Pending Registers.	

Table 16.10 Centronics Reset Interrupt Pending Field.

Centronics Reset Interrupt Enable

Value	Action
'1'	Signal Pending Reset Interrupts to the interrupt controller.
'0'	Do not signal the interrupt controller (default).

Table 16.11 Centronics Negotiation Interrupt Enable Field.

Printer Error ('PError') Field:

Value	Meaning
'1'	Printer Error
'0'	Normal (default).

Table 16.12 Printer Error Field ('PError') Field Encoding.

Printer On Line Select ('Select') Field:

Value	Meaning
'1'	Printer on line
'0'	Printer off line (default).

Table 16.13 Select ('Select') Field Encoding.

Printer Fault ('Fault') Field:

Value	Meaning
'1'	Printer Fault
'0'	Normal (default).

Table 16.14 Printer Fault ('Fault') Field Encoding.

Printer Acknowledge Negated ('AckN') Field:

Value	Meaning
'1'	Normal
'0'	Printer Acknowledge (default).

Table 16.15 Printer Acknowledge Negated ('AckN') Field Encoding.

Printer Busy ('Busy') Field:

Value	Meaning
'1'	Printer Busy
'0'	Printer not Busy (default).

Table 16.16 Printer Busy Field ('Busy') Field Encoding.

Centronics Control Register ('CentControlReg')

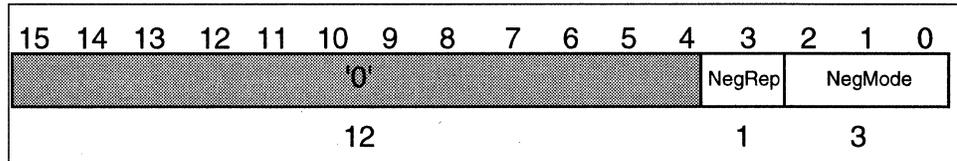


Figure 16.4 Centronics Control Register ('CentControlReg').

Figure 16.4 illustrates the Centronics Control Register and its fields. Additional programming instructions are located in Table 16.17, Table 16.18, and Table 16.19.

Bit	Description
3	Negotiation XFlag Reply (NegRep)
2:0	Negotiation Mode (NegMode)

Table 16.17 Centronics Control Register ('CentControl') Bit Assignments.

Negotiation XFlag Reply ('NegRep') Field:

Value	Action
'1'	Nibble Mode: mode requested supported. Other Modes: mode requested not supported.
'0'	Nibble Mode: mode requested not supported. Other Modes: mode requested supported (default).

Table 16.18 Negotiation XFlag Reply ('NegRep') Field Encoding.

Negotiation Mode ('NegMode') Field:

Value	Action
b'111'	Termination
b'110'	Extensibility Link
b'101'	CPU Control (reserved)
b'100'	EPP
b'011'	ECP
b'010'	Byte
b'001'	Nibble
b'000'	Compatible (default)

Table 16.19 Negotiation Mode ('NMode') Field Encoding.

Centronics Nibble Data Register ('CentNibbleDataReg')

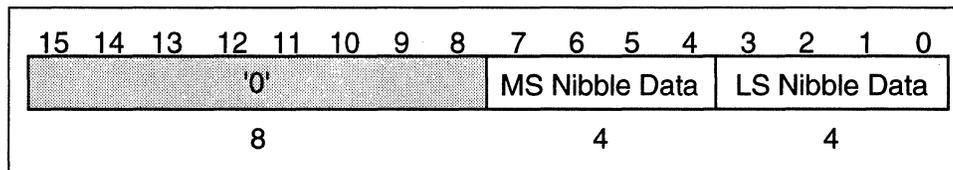


Figure 16.5 Centronics Nibble Data Register ('CentNibbleDataReg').

Bit	Description
7:4	Most Significant Nibble Data
3:0	Least Significant Nibble Data

Table 16.20 Centronics Nibble Data Register ('CentNibbleDataReg') Bit Assignments.

Figure 16.5 and Figure 16.6 are illustrations of the Centronics Nibble Data and Host Status Registers. Additional programming instructions for these registers are located in Table 16.20, Table 16.21, Table 16.22, Table 16.23, Table 16.24, Table 16.25.

Centronics Host Status Register ('CentHostStatusReg')

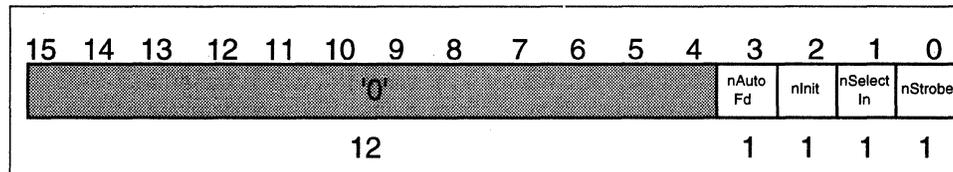


Figure 16.6 Centronics Host Status Register ('CentHostStatusReg').

Bit	Description
3	AutoFeed Negated (nAutoFd)
2	Initialize Negated (nInit)
1	Select In Negated (nSelectIn)
0	Host Strobe Negated (nStrobe)

Table 16.21 Centronics Host Status Register ('CentHostStatusReg') Bit Assignments.

AutoFeed Negated ('nAutoFeed') Field:

Value	Meaning
'1'	Normal (default).
'0'	AutoFeed mode.

Table 16.22 AutoFeed Negated ('nAutoFeed') Field Encoding.

Initialize Negated ('nInit') Field:

Value	Meaning
'1'	Normal (default).
'0'	Initialize the printer.

Table 16.23 Initialize Negated ('nInit') Field Encoding.

Select In Negated ('nSelectIn') Field:

Value	Meaning
'1'	Don't select this printer (default).
'0'	Select this printer.

Table 16.24 Select In Negated ('nSelectIn') Field Encoding.

Host Strobe Negated ('nStrobe') Field:

Value	Meaning
'1'	Host Strobe pulse de-asserted (default).
'0'	Host Strobe pulse asserted.

Table 16.25 Host Strobe Negated ('nStrobe') Field Encoding.

Centronics Minimum Delay Register ('CentDelayReg')

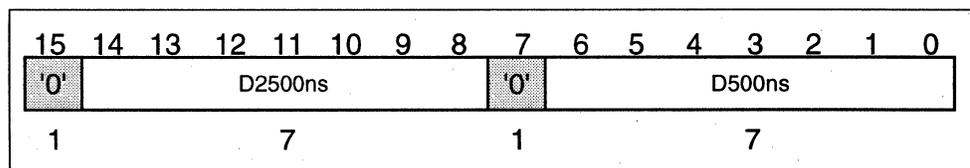


Figure 16.7 Centronics Minimum Delay Register ('CentDelayReg').

Figure 16.7 is an illustration of the Centronics Minimum Delay Register. Refer to Table 16.26, Table 16.27, Table 16.28, and Table 16.29 for additional programming information.

Bit	Description
15	Reserved. Must be written as '0'.
14:8	2500ns Delay Type Field (D2500ns)
7	Reserved. Must be written as '0'.
6:0	500ns Delay Type Field (D500ns)

Table 16.26 Centronics Minimum Delay Register ('CentDelayReg') Bit Assignments.

2500ns Delay Type Field ('D2500ns') Field:

Value	Action
'0x7f'	0x7f clock delay.
...	
'0x01'	0x01 clock delay.
'0x00'	undefined. (default).

Table 16.27 2500ns Delay Type Field ('D2500ns') Field Encoding.

500ns Delay Type Field ('D500ns') Field:

Value	Action
'0x7f'	0x7f clock delay.
...	
0x01'	0x01 clock delay.
'0x00'	undefined. (default).

Table 16.28 500ns Delay Type Field ('D500ns') Field Encoding.

CPU MHz	2500ns	500ns
33	0x53	0x11
25	0x3f	0x0d
20	0x32	0x0a
16	0x28	0x08

Table 16.29 Example Settings for Delay Type Fields.

Timing Diagram

Figure 16.8 illustrates a typically classic compatible mode transaction.

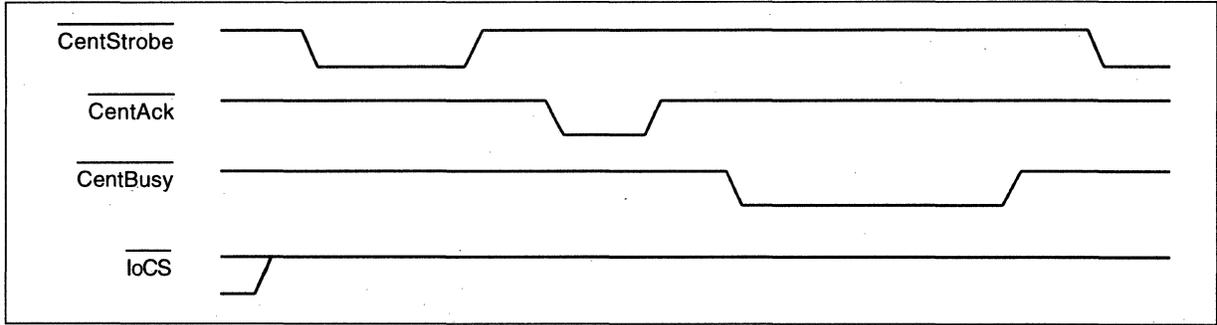


Figure 16.8 Typically Classic Compatible Mode Transaction.

System Example

Figure 16.9 illustrates typical parallel port system connections.

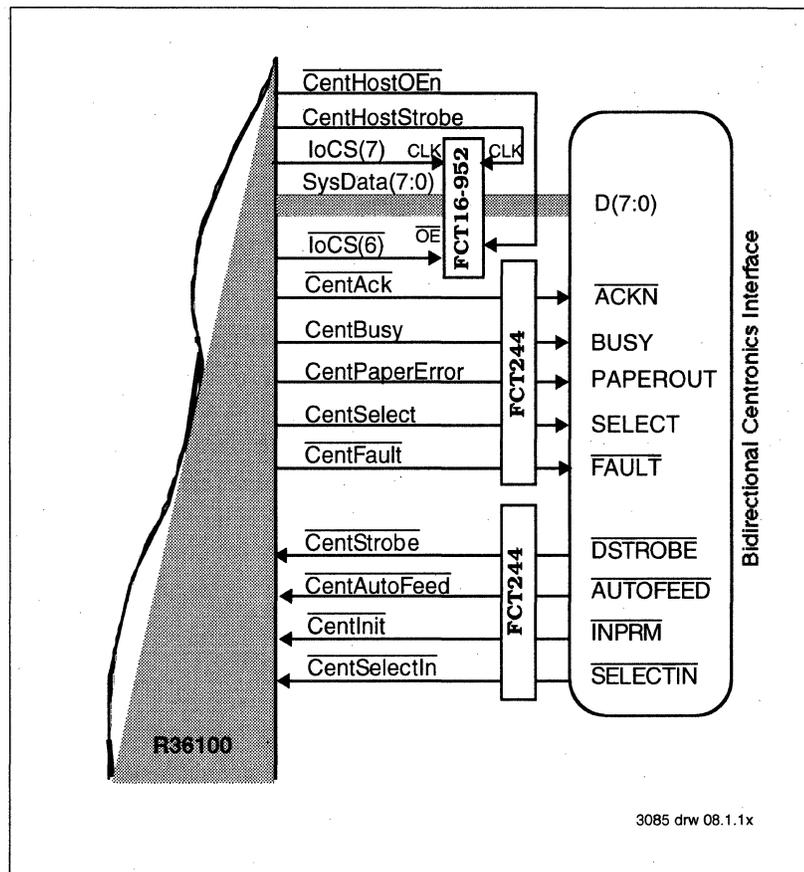


Figure 16.9 Typical Parallel Port System Connections.

Note: The virtual address for the external Centronics data register (in this example, FCT16-952) must have address bit 15 low.



Integrated Device Technology, Inc.

Introduction

The IDT R36100 RISController integrates bus controllers and peripherals around the R30xx family CPU core. One of the many on-chip peripherals is the *Laser Printer Video Port* as described in this chapter.

This chapter will provide an overview of the Laser Printer Video Port register interface, a description of the signal pins, and various aspects of the signal timing. Figure 17.1 shows a block diagram of the Laser Printer Video Port

Features

- 1-bit serial stream laser printer or raster engine interface
- Supports 4-pin engine interface standard
- On-chip 4-word Transmit FIFO
- Programmable margin widths and page length
- DMA auto-initiate via internal interrupt

Block Diagram

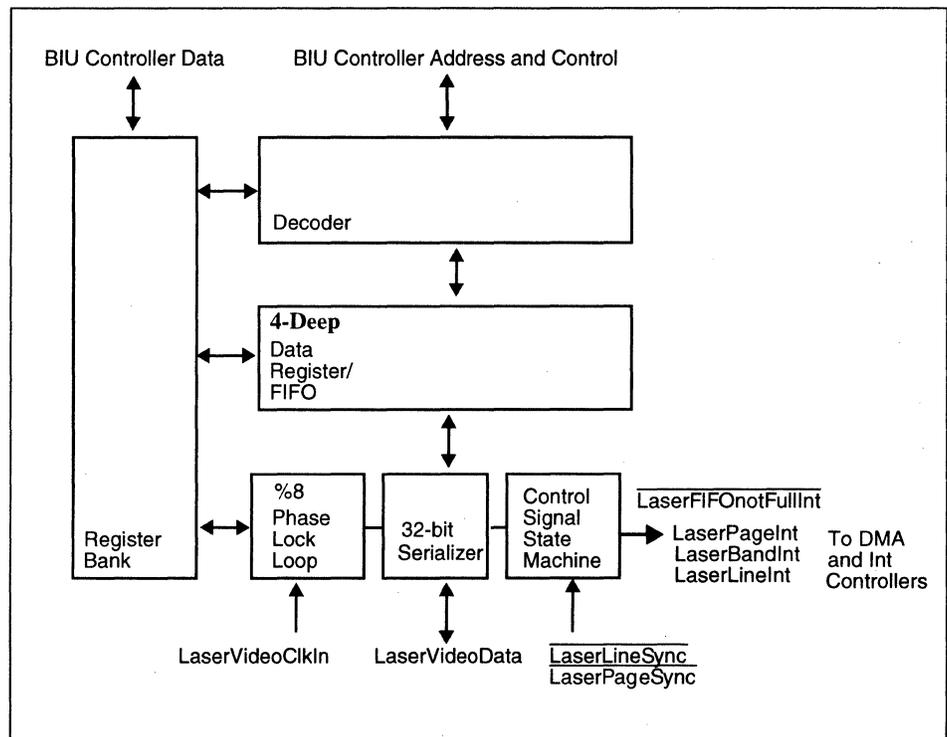


Figure 17.1 Block Diagram of the Laser Printer Video Port

Overview

The Laser controller's overall purpose is to take a laser printer video image stored in DRAM frame buffer memory, serialize it, and send it on to the laser printer imaging drum engine.

When the R36100 is used as a laser printer controller, the CPU receives Postscript or PCL page description language (PDL) commands from the serial or the parallel port. The software (PDL interpreter) translates the PDL commands into an imaged video bit-map and stores the bit-map into a frame buffer in DRAM. Once the image is ready, the CPU uses the on-chip Laser Controller to interface to the mechanics of the laser printer engine. The interface is used to let the engine mechanically load a page. The engine then begins generating page sync (LaserPageSync) and line sync (LaserLineSync) signals.

Once these signals begin, the Laser Controller must give a continuous serial bit stream from the data in its serializer buffer. The serializer buffer is supplied by a 4-word transmit FIFO which in turn must be updated from the frame buffer DRAM. With DMA, data can autonomously and automatically be loaded from DRAM into the transmit FIFO. Meanwhile, the CPU continues to finish up imaging the page or to begin receiving and imaging the next page. The primary data path interface of the Laser controller consists of a 4-word transmit FIFO. This FIFO can be updated by either CPU single word writes or, more efficiently, by a DMA channel.

The 32-bit words of video image data are moved through the FIFO to a separate 32-bit serializer buffer with the Most Significant Bit (MSB) being the 1st bit shifted out and then driven from the CPU to the external Laser engine. Optionally, by using the 'Dir' field, the serializer can shift in the opposite direction. The video is a one bit wide serial stream of digital data. For a Canon type printer engine where the video clock is generated by the controller, there is a phase lock loop (PLL) on the Laser Controller. The PLL synchronizes the video data to the assertion of LaserLineSync. Six clocks after LaserLineSync is de-asserted, the Video Data begins shifting out.

When the PLL option is selected, the input clock for the video is at eight times the true dot frequency. For Laser engines that supply the video clock, there is a PLL bypass option where the video clock is the same frequency as supplied. The Laser Controller serializer can shift data in either the left or right direction to support engines with duplex printing capabilities, which allows printing on both sides of the paper.

Vertical and Horizontal page margins are supported. The vertical skip counter tells the Laser controller how many lines to skip before starting to output video data. The vertical skip counter starts counting on LaserPageSync, and decreases every LaserLineSync including the initial LaserLineSync. The horizontal skip counter sets the horizontal page margin by telling how many dots to skip from the beginning of each line before starting to print. The horizontal counter is loaded and starts counting on LaserLineSync and decrements every video clock. After finishing the skip time, the actual data in the serializer begins to be shifted out. At the end of a line (as determined by the Line Pixel Count), any bits remaining in the serializer buffer are deleted/flushed and the next FIFO word is loaded.

FIFO Depth:

When the FIFO is not full, it will generate a not full status interrupt to notify the DMA controller, or CPU, that it can now be refilled.

Internal PLL frequency:

The internal PLL frequency is capable of supporting a variety of engines:

- 2.5MHz %8, -> 40MHz PLL (e.g. 300DPI)
- 10MHz %8 -> 80MHz PLL (e.g. 600DPI)
- 15MHz %8 -> 100-120MHz PLL (e.g. high performance 600DPI)

Reset Restriction:

For proper operation, $\overline{\text{SysReset}}$ width must be longer than 5 internally derived Video Clock cycles.

Interrupt Descriptions **$\overline{\text{LaserFIFONotFullInt}}$**

Asserts when the Laser Video FIFO has one or more empty entries. Stays asserted until the FIFO is full.

 $\overline{\text{LaserPageInt}}$

Asserts for 1 clock pulse at the leading edge of $\overline{\text{LaserPageSync}}$.

 $\overline{\text{LaserBandInt}}$

Asserts for 1 clock pulse on the second bit of the last line of band.

 $\overline{\text{LaserLineInt}}$

Asserts for 1 clock pulse when the first bit of the 1 word is left at the end of a line and indicates that the Laser Video FIFO is about to become empty.

Pin Descriptions**Laser Printer Video Control Interface Signals**

Any one of the CPU I/O port chip selects, $\overline{\text{IoCS}}()$ can be used to support the byte wide EngineStrobe data to the engine control and the byte wide data from Engine Output status to the R36100. Note that the Laser Video pins are physically multiplexed with the PIO pins and must be appropriately programmed before use.

Laser Printer Video Interface Signals

The five Laser video signal pins are multiplexed with PIO pins, and thus each of those PIO pins must be programmed to use the pins as Laser video pins.

LaserVideoClkIn **Input**

The video clock is generated as a pixel clock which strobes the Laser-Video Data stream. The video clock is the primary clock for controlling the pixel rate of the video engine interface and is derived either directly from LaserVClkIn or indirectly from an internal PLL. LaserVideoClkIn when used with the PLL is 8x the actual pixel rate. LaserVideoClkIn when not used with the PLL is the same as the actual pixel rate. The maximum internally derived video clock frequency is 120MHz.

Note: PIO(25) must be programmed appropriately for LaserVideoClkIn to be fully functional.

LaserVideoData **Output**

Laser Printer Video Data: Serial data stream connected to the print head.

Note: PIO(8) must be programmed appropriately for LaserVideoClkIn to be fully functional.

LaserLineSync **Input**

Laser Printer Line Sync: Input from the engine print head that indicates that the next line of data is to begin the transfer. The LaserVideoData stream begins to output data (assuming that the horizontal margin is 0) 6 clocks after the de-assertion of the $\overline{\text{LaserLineSync}}$ pulse. The $\overline{\text{LaserLineSync}}$ pulse width must be at least as wide as the internally derived video clock period.

Note: PIO(24) must be programmed appropriately for LaserVideoClkIn to be fully functional.

LaserPageSync **Input**

Laser Printer Page Synchronization: Input from the engine print head indicating that the next page is to begin its transfer. The $\overline{\text{LaserPageSync}}$ pulse width must be at least as wide as the internally derived video clock period.

Note: PIO(3) must be programmed appropriately for LaserVideoClkIn to be fully functional.

Register Definitions

Table 17.1 lists the addresses and descriptions for the Laser Printer Video Controller Registers. Note that Big Endian software must offset these addresses by b'10 (0x2), except for the Video Data Register, which is a full 32-bit wide register.

Phys. Address	Description
0xFFFF_ED00	Video Control Register
0xFFFF_ED04	Video Vertical Skip Count Register
0xFFFF_ED08	Video Horizontal Skip Count Register
0xFFFF_ED0C	Video Band Line Count Register
0xFFFF_ED10	Video Pixel Word Count Register
0xFFFF_ED20	Video Data Register

Table 17.1 Laser Printer Video Controller Registers.

Laser Video Control Register

('LaserControlReg')

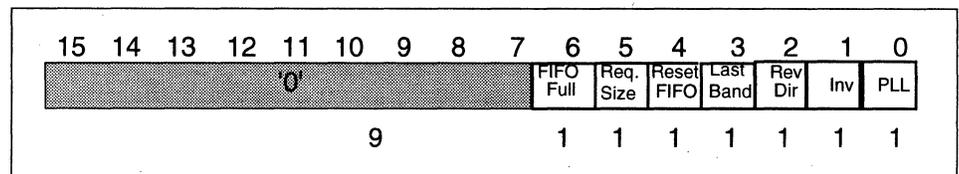


Figure 17.2 Laser Video Control Register ('LaserControlReg').

Figure 17.2 illustrates the Laser Video Control Register fields. Additional programming information for these fields is located in Table 17.2, Table 17.3, Table 17.4, Table 17.5, Table 17.6, Table 17.7, Table 17.8, and Table 17.9.

Bit	Description
6	Video FIFO is Full
5	Request Size
4	Reset FIFO
3	Last Band
2	Reverse Video Stream Direction
1	Inverse Video
0	PLL On

Table 17.2 Laser Video Control Register ('LaserControlReg') Bit Assignments.

Video FIFO Full

Value	Meaning
'1'	Video FIFO is full
'0'	Video FIFO is not Full (default)

Table 17.3 FIFO Full Field.

Request Size ('ReqSize')

Request size affects the Laser FIFONotFullInt flag as to how many entries are used for the flag.

Value	Action
'1'	Reserved.
'0'	Request data when 1 or more slots are empty (default).

Table 17.4 Request Size Field.

Reset FIFO ('ResetFifo')

A FIFO reset requires the following two steps:

1. Writing a '1.' to reset the FIFO.
2. Writing a '0.' to 'un' reset the FIFO.

Value	Action
'1'	When written with '1', clears video FIFO to empty.
'0'	No action (default).

Table 17.5 Reset FIFO Field.

Last Band ('Last') Field:

Last Band Field must be set between pages before a new page sync.

Value	Action
'1'	Enable Interrupt at the end of this Last Band.
'0'	No action (default).

Table 17.6 Last Band ('Last') Field Encoding.

Video Direction ('Dir') Field:

Value	Action
'1'	Video Stream shifts from MSB to LSB (normal direction).
'0'	Video Stream shifts from LSB to MSB (default) (reverse direction).
Note: In the initial version of the R36100, the DMA controller does not support the reverse direction; software is responsible for providing data in this order.	

Table 17.7 Video Direction ('Dir') Field Encoding.

Inverse Video Data ('Inv') Field:

Value	Action
'1'	Video Stream data is inverted.
'0'	Video Stream data is normal (default).

Table 17.8 Inverse Video Data ('Inv') Field Encoding.

Phase Lock Loop Enabled ('PLL') Field:

Value	Action
'1'	PLL enabled (video clock = 8x PLL).
'0'	PLL disabled (video clock = LaserVClkIn (default)).

Table 17.9 Phase Lock Loop Enabled ('PLL') Field Encoding.

Laser Video Vertical Skip Register ('LaserVertSkipReg')

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
VS	Dis	0	Vertical Skip Count -1												
1	2												13		

Figure 17.3 Laser Vertical Skip Register ('LaserVertSkipReg')

The Laser Vertical Skip Register, shown in Figure 17.3, must only be altered when the controller is not in the process of vertical skipping.

Additional programming information for this register is located in Table 17.10, Table 17.11, and Table 17.12.

Bit	Description
15	Vertical Skip Disable ('VSDis')
12:0	Vertical Skip Count -1 ('VSCount')

Table 17.10 Laser Video Vertical Skip Register ('LaserVertSkipReg') Bit Assignments.

Laser Vertical Skip Dis ('VSDis') Field:

The VSDis Field is used to program the '0' skip case.

Value	Action
'1'	Vertical Skip is disabled (default).
'0'	Vertical Skip Count is enabled starting with the next page.

Table 17.11 Vertical Skip Enable ('VSDis') Field Encoding.

Vertical Skip Count ('VSCount') Field:

Note the actual down counter is internal and is not accessible by the programmer. This Count/Compare register is loaded into the internal count register at first VideoClkIn after SysReset and at page sync.

Value	Action
'0x7FF'	Number of Lines to skip at the beginning of a page 1.
...	
'0x000'	Skip 1 line (default).

Table 17.12 Vertical Skip Count ('VSCount') Field Encoding.

Laser Video Horizontal Skip Register ('LaserHorSkipReg')

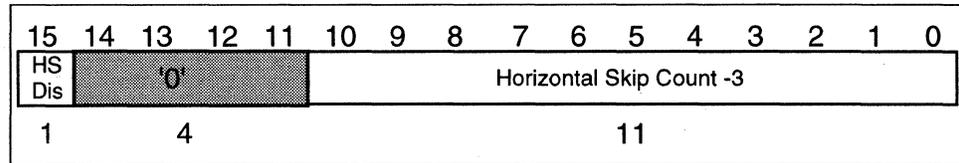


Figure 17.4 Laser Video Horizontal Skip Register ('LaserHorSkipReg').

The Laser Video Horizontal Skip Register, shown in Figure 17.4, must only be altered when the controller is not in the process of horizontal skipping, preferably between pages. Note the actual down counter is internal and is not accessible by the programmer. This Count/Compare register is loaded into the internal count register during line sync.

Additional programming information for the fields of this register is located in Table 17.13, Table 17.14, and Table 17.15.

Bit	Description
15	Horizontal Skip Disable ('HSDis')
10:0	Horizontal Skip Count - 4 ('HSCount')

Table 17.13 Laser Video Horizontal Skip Register ('LaserHorSkipReg') Bit Assignments.

Horizontal Skip Enable ('HSDisEn') Field:

Note: Whether or not the HSD Field is on or off, five clocks are skipped after LaserLineSync de-asserts.

Value	Action
'1'	Horizontal Skip Count is disabled (default).
'0'	Horizontal Skip is enabled for the next line.

Table 17.14 Horizontal Skip Enable ('HSDisEn') Field Encoding.

Horizontal Skip Count ('HSCount') Field:

Note the actual down counter is internal and is not accessible by the programmer. This Count/Compare register is loaded into the internal count register during line sync. Therefore, 5+ HSCount +4 pixels are skipped.

Value	Action
'0x7FF'	Number of pixels +4 to skip at the beginning of a line.
...	
'0x000'	Skip 1 pixel (default).

Table 17.15 Horizontal Skip Count ('HSCount') Field Encoding.

Laser Video Band Line Count Register ('LaserBandLineCountReg')

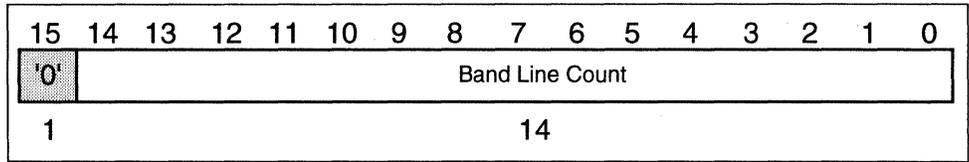


Figure 17.5 Laser Video Band Line Count Register ('LaserBandLineCountReg').

Bit	Description
13:0	Vertical Band Line Count ('BandLineCount')

Table 17.16 Laser Video Band Line Count Register ('LaserBandLineCountReg') Bit Assignments.

Band Line (Lines per Band/Page) Count ('BandLineCount') Field:

The Laser Video Band Line Count Register is shown in Figure 17.5. Additional programming information for the fields of this register is given in Table 17.16 and in Table 17.17.

Note: The actual down counter is internal and is not accessible by the programmer. This Count/Compare register is loaded into the internal count register during page sync and at the end of each band. Typically, there are several bands per page.

Value	Action
'0x3FFF'	Number of lines per band/page (not including vertical skip count) - 1.
...	...
'0x0000'	1 line per band/page (not including vertical skip count) (default).

Table 17.17 Band Line (Lines per Band/Page) Count ('BLCount') Field Encoding.

Laser Video Horizontal Count Register ('LaserPixelWordCountReg')

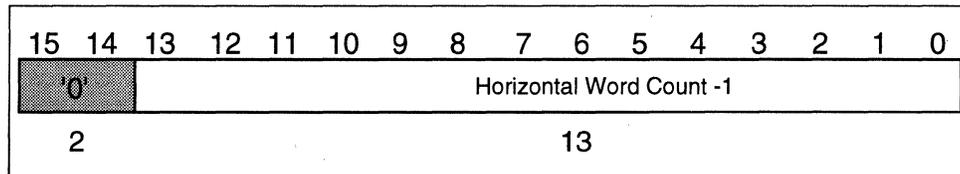


Figure 17.6 Laser Video Horizontal Word Count Register ('LaserHorizPixelCountReg').

Bit	Description
13:0	Video Horizontal Pixel Word Count - 1

Table 17.18 Laser Video Horizontal Pixel Word Count Register ('LaserPixelWordCountReg') Bit Assignments.

Figure 17.6 illustrates the fields of the Laser Video Horizontal Word Count Register. Additional programming information for these fields is located in Table 17.18 Programming instructions for The Video Vertical Pixel Count Field is located in Table 17.19.

Video Pixel Word Count ('PixelWordCount') Field:

Note: The actual down counter is internal and is not accessible by the programmer. This Count/Compare register is loaded into the internal count register during line sync. If the PixelWordCount field is not divisible by 32, then the remaining bits are unidentified.

Value	Action
'0x3FFF'	Number of words (32 pixels/word) per Line (not including Skip Count).
...	...
'0x0000'	1 word per line (not including Skip Count) (default).

Table 17.19 Video Vertical Pixel Count ('PixCount') Field Encoding.

Laser Video Data Register ('LaserDataReg')

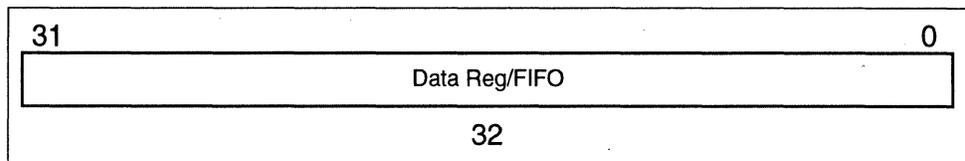


Figure 17.7 Laser Video Data Register ('LaserWordDataReg').

The Data Register, shown in Figure 17.7, is written to with a store word "sw" instruction (or the DMA equivalent). Figure 17.8 provides a R36100 Video Interface System example.

System Example

Figure 17.8 is an example of an R36100 video interface system.

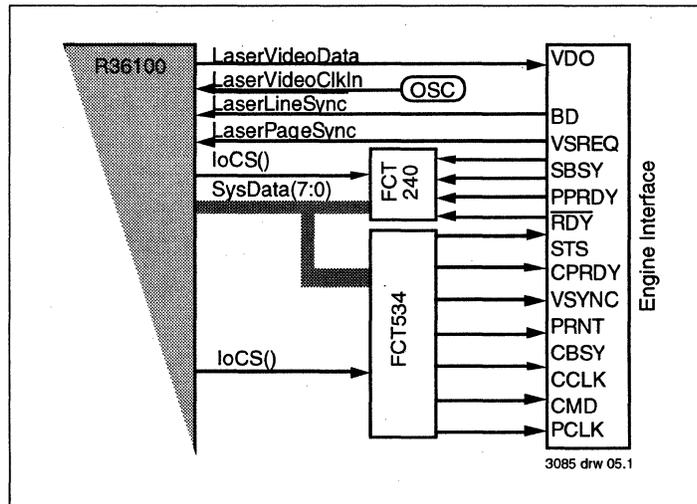


Figure 17.8 R36100 Video Interface System Example



Introduction

This chapter discusses the reset initialization sequence required by the R36100. Also included is a discussion of the configuration mode selectable features of the processor and of the software requirements of the boot program.

There are a number of selectable features in the R36100. These mode selectable features are determined by the polarity of the appropriate reset configuration mode inputs when the rising edge of $\overline{\text{SysReset}}$ occurs.

Reset Timing

The R36100 requires a very simple reset sequence. There are only two concerns for the system designer:

- That the set-up time and hold requirements of the reset configuration mode feature inputs with respect to the rising edge of $\overline{\text{SysReset}}$ are met.
- That the minimum $\overline{\text{SysReset}}$ pulse width is satisfied.

Reset Configuration Mode Features

The R36100 has features that are determined at reset time. This is achieved by using a latch internal to the CPU; this latch samples the contents of the reset mode feature bus at the negating edge of $\overline{\text{SysReset}}$. The encoding of the mode selectable features on the reset mode feature bus is described in Table 18.1.

Pin	Mode Feature
ExcSInt(0)	BigEndian/LittleEndian
ExcSInt(1)	BootProm8
ExcSInt(2)	BootProm16

Table 18.1 R36100 Reset Configuration Mode Features

Reset Configuration Mode Pin Descriptions

Exception Signals

SysReset Input

System Reset is a master processor reset active-low-input signal that initializes the processor. The processor's optional features are established during the last cycle of reset, using the reset configuration mode inputs from EXCSInt(2:0).

Internally to the chip, $\overline{\text{SysReset}}$ is further hardened (relative to the R36100) to accept slow rise times by the use of hysteresis and the elimination of the high pass TTL-level filtering.

ExcSInt(2:0) Input

Exception Synchronized Interrupt: These signals are the same as the R3051 $\overline{\text{SInt}}(2:0)$ signals except for the Reset Configuration Modes.

Value	Description
'11'	32-bit wide (non-interleaved) Boot Prom
'10'	8-bit wide Boot Prom
'01'	16-bit wide Boot Prom
'00'	64-bit wide (interleaved 32-bit wide) Boot Prom

Table 18.2 Boot Prom Reset Configuration Modes for ExcSIntN(2:1) pins.

Note: The values of the Reset Initialization Vector for the Boot PROM are inverted relative to the internal size field in the configuration register of the Memory Controller.

BigEndian

Use Big Endian Addressing: if asserted (active high), the processor will operate as a big-endian machine, and the RE bit of the status register would then allow little-endian tasks to operate in a big-endian system. If negated (inactive low), the processor will operate as a little-endian machine, and the RE bit will allow big-endian tasks to operate on a little-endian machine.

BootProm8

8-bit Boot PROM Mode. If asserted (active low), this mode will cause the port size mapping register to initialize all memory sub-regions to 8-bit ports instead of 32-bit ports. Thus an 8-bit boot PROM can be used to initialize the R3041. This mode can only be asserted if $\overline{\text{BootProm16}}$ is de-asserted. Table 18.2 shows the encoding of this bit at reset.

BootProm16

16-bit Boot PROM Mode: if asserted (active low), this mode will cause the port size mapping register to initialize all memory sub-regions to 16-bit ports instead of 32-bit ports. A 16-bit boot PROM can be used to initialize the R3041. This mode can only be asserted if $\overline{\text{BootProm8}}$ is de-asserted. Table 18.2 shows the encoding of this bit at reset.

R3000A Equivalent Modes

The R3000A features a number of modes, which are selected at Reset time. Although most of those modes are irrelevant, a number of equivalences can be made:

- IBlkSize = 4 word refill.
- DBlkSize = 1 or 4 word refill, depending on the DBlockRefill mode as selected in the CP0 Cache Configuration register.
- Reverse Endianness capability enabled.
- Instruction Streaming enabled.
- Partial Word Stores enabled.

Other modes of the R3000A pertain primarily to its cache interface, which is incorporated within the R36100 and transparent to users of this processor.

Reset Behavior

While Reset is asserted, the processor maintains its interface in a state that allows the rest of the system to also be reset. Specifically:

- SysClk operates at one-half the ClkIn frequency.
- SysData() is tri-stated
- SysALEn is driven de-asserted (high).
- Control signals are driven de-asserted (high).
- SysAddr() and SysDiag functions are driven (value undefined).

The R36100 samples for the negation of $\overline{\text{SysReset}}$ relative to a falling edge of $\overline{\text{SysClk}}$. On a rising edge of $\overline{\text{SysClk}}$, 6 cycles after the negation of $\overline{\text{SysReset}}$ is detected, the processor initiates a read request for the instruction located at the Reset Exception Address Vector. These cycles are a result of:

- $\overline{\text{SysReset}}$ input synchronization performed by the CPU. The $\overline{\text{SysReset}}$ input uses special synchronization logic, thus allowing $\overline{\text{SysReset}}$ to be negated asynchronously to the processor. This synchronization logic introduces a two cycle delay between the external negation of $\overline{\text{SysReset}}$ and the negation of $\overline{\text{SysReset}}$ to the execution core.
- Internal clock cycles in which the execution core flushes its pipeline, before it attempts to read the exception vector.
- One additional cycle for the read request to propagate from the internal execution core to the read interface, as described in Chapter 8.

Boot Software Requirements

Basic mode selection is performed using hardware during the reset sequence, as discussed in the mode initialization section. However, there are certain aspects of the boot sequence that must be performed by software.

The assertion and subsequent negation of reset forces the CPU to begin execution at the reset vector, which is physical address 0x1FC0_0000. This address resides in uncached, un-mapped memory, and thus does not require that the caches be initialized for the processor to execute boot code.

The processor must perform the following activities during boot:

- Initialize the CPO Status Register. The processor must be assured of having the kernel enabled to perform the boot sequence. Typically, a `mtc0 rx, CO_SR` instruction is one of the first few instructions in the boot sequence. Specifically, co-processor usable bits, and cache control bits, must be set to the desired value before any data references (cached or uncached), diagnostics or initialization occur.
- Initialize the CPO Configuration Registers. The software should decide on the Cache Configuration, Port Sizes, and Bus Control during initialization.
- Initialize the caches. The processor must determine the sizes of the on-chip caches, and flush each entry, as discussed in Chapter 3. This must be done before the processor attempts to execute cacheable code.
- Re-initialize CPO Registers. The processor should establish appropriate values in various CPO registers, including:
 - The IM bits of the status register.
 - The BEV bit.
 - Initialize KUp/IEp so that user state can be entered using a RFE instruction.
- Initialize on-chip memory and I/O controllers. The boot software should establish the appropriate timing parameters, control options, timer values, PIO uses, etc., as appropriate to the particular system.
- Enter User State.

Branch to the first user task and perform an RFE to enter the user mode.

Detailed Reset Timing Diagrams

The timing requirements of the processor reset sequence are illustrated below. The timing diagrams reference AC parameters whose values are contained in the R36100 data sheet.

Reset Pulse Width

There are two parameters to be concerned with: the power on reset pulse width and the warm reset pulse width. The minimum number of clock cycles for a warm-reset depends on whether the system relies on the internal pull-ups of the mode-vectors, or if they are actively driven. If using the internal pull-ups, a considerably longer time must be allowed, since the pull-up values are rather weak.

Figure 18.1 illustrates the power on reset requirements of the R36100 family. And Figure 18.2 illustrates the warm reset requirements of the processor.

Mode Initialization Timing Requirements

The mode initialization vectors are sampled by an internal transparent latch, whose output enable is directly controlled by the $\overline{\text{SysReset}}$ input of the processor. The internal structure of the processor is illustrated in Figure 18.3. As illustrated in Figure 18.4, the mode vectors have a set-up and hold time with respect to the rising edge of $\overline{\text{SysReset}}$.

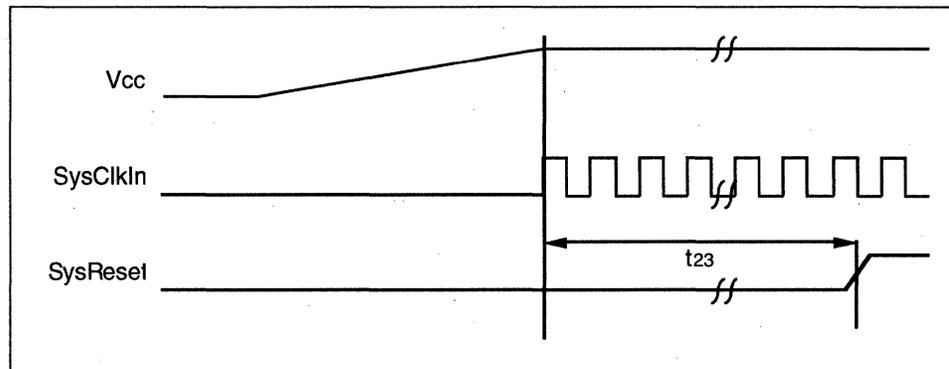


Figure 18.1 Cold Start

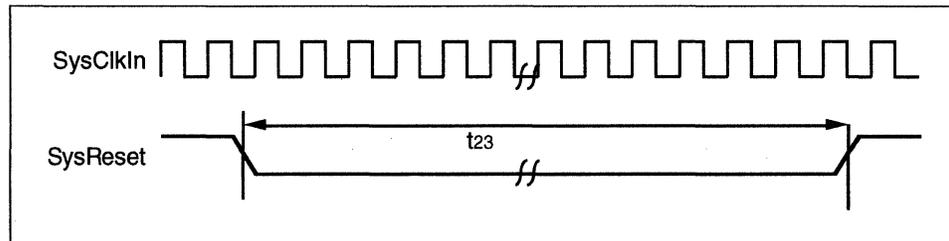


Figure 18.2 Warm Reset

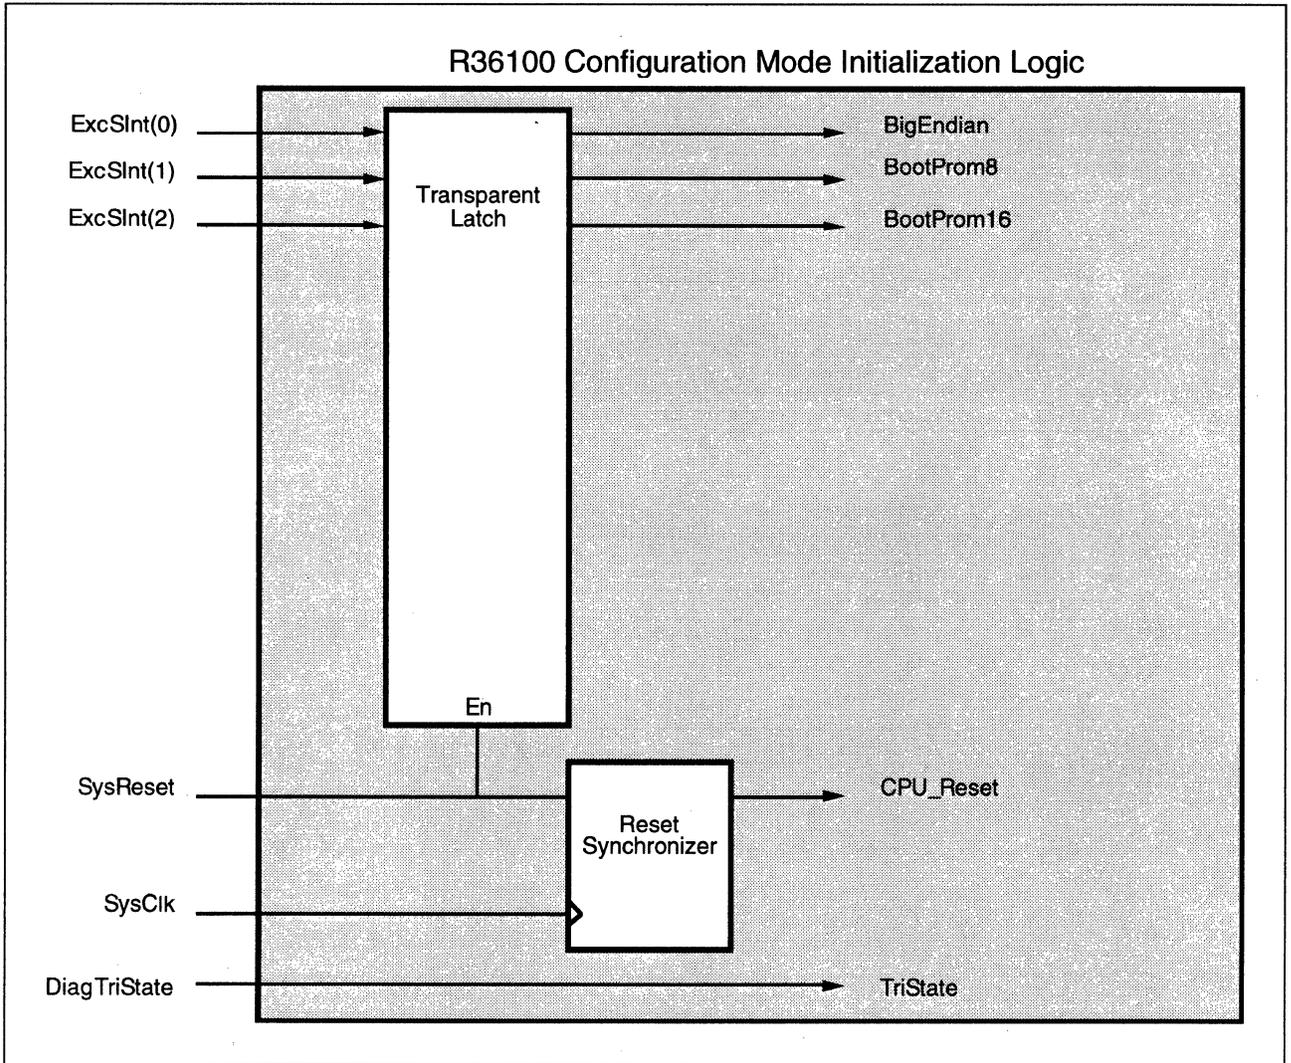


Figure 18.3 Configuration Mode Initialization Logic

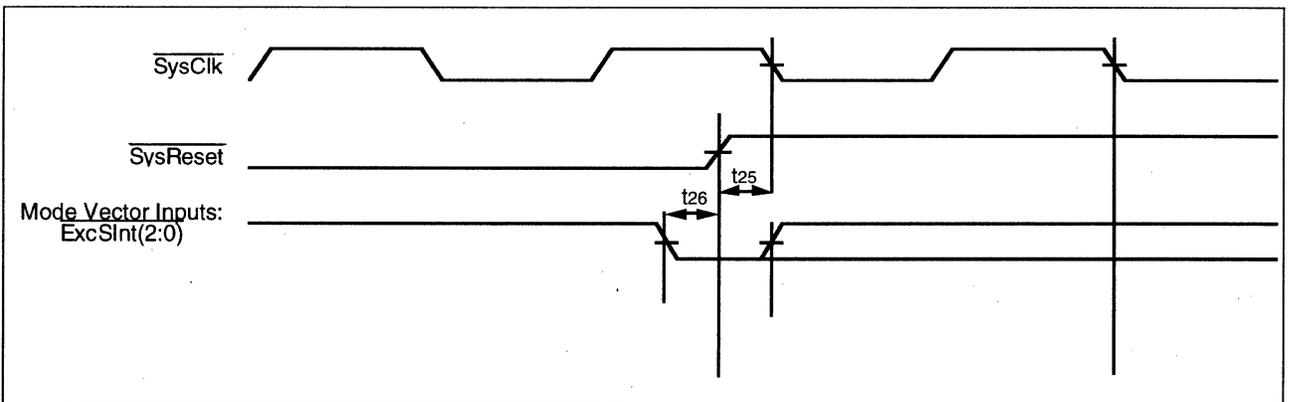


Figure 18.4 Mode Vector Timing

Reset Setup Time Requirements

The reset signal incorporates special synchronization logic that allows it to be driven from an asynchronous source. This allows the processor $\overline{\text{SysReset}}$ signal to be derived from a simple circuit, such as an RC network, with a time constant long enough to guarantee that the reset pulse width requirement is met.

Such a system should buffer the RC circuit such that a sufficiently fast monotonic rise time is generated which is capable of synchronously resetting any external state machines and logic at the same time as resetting the CPU.

The $\overline{\text{SysReset}}$ set-up time parameter can be thought of as the amount of time $\overline{\text{SysReset}}$ must be negated before the rising edge of SysClk , for guaranteed recognition. Failure to meet this requirement will delay the internal recognition of the end of reset by one clock cycle. This does not affect the timing of the sampling of the mode initialization vectors. Figure 18.5 illustrates the set-up time parameter of the R36100.

ClkIn Requirements

The input clock timing requirements are illustrated in Figure 18.6. The system designer does not need to be explicitly aware of the timing relationship between ClkIn and $\overline{\text{SysClk}}$. Note that $\overline{\text{SysClk}}$ is driven even during the $\overline{\text{SysReset}}$ period as long as ClkIn is provided.

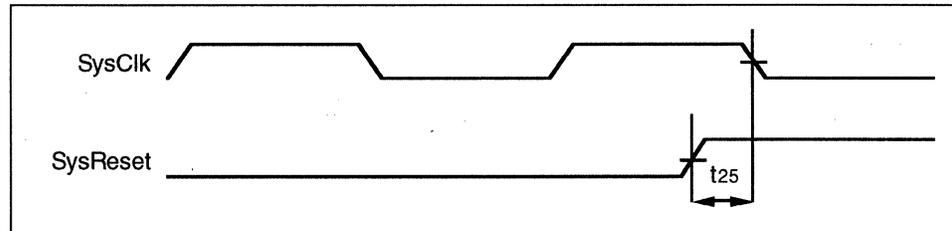


Figure 18.5 Reset Timing

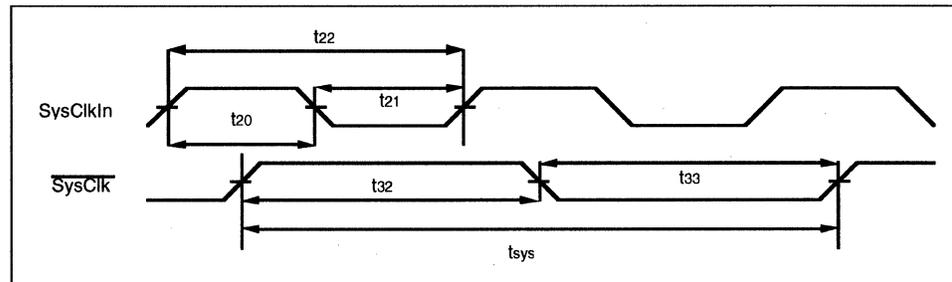


Figure 18.6 R36100 Clocking



Integrated Device Technology, Inc.

Introduction

This chapter discusses features included to facilitate debugging of R36100-based systems. These features are intended to be used by an in-circuit emulator, in-circuit tester, board level tester, logic analyzer, hardware modeler, or similar tool.

Features

- Hardware trace/halt support, cache write suppression, and KO preservation
- Cause register write option of the exception code bits (CP0)
- Instruction stepping support via virtual address debug trace watch register
- The ability of the processor to have instruction and data cache misses forced, thus allowing all internal cache accesses to be displayed on the bus interface.
- The ability to tri-state all output pins including $\overline{\text{SysClk}}$, thus allowing an in-circuit emulator or tester to drive and control the output pins directly.
- The ability to deterministically set the phase relationship of the $\overline{\text{SysClk}}$ output relative to the ClkIn input. This feature allows board level testers and hardware modelers to control the $\overline{\text{SysClk}}$ output.
- The ability to distinguish data and instruction accesses, allowing logic analyzers to do instruction disassembly.
- A software breakpoint instruction.

Note: The features described in this chapter are intended for initial debug or production testing rather than for use in an end-user system.

The following are several debug/emulator hooks included in the normal functioning chip:

- tri-state-able outputs
- tracepoint register
- extended CP0 cache configuration register
- JTAG output scan path

Tri-State-able Outputs

The tri-state-able outputs feature uses a dedicated input pin that if asserted causes all outputs (including $\overline{\text{SysClk}}$) on the chip to tri-state. This feature is used in in-circuit manufacturing tests and by in-circuit emulators with non-socketed target CPUs.

Tracepoint Registers

The tracepoint registers consist of two memory mapped virtual address registers and a control register. When enabled through the control register, the tracepoint registers cause an exception when the virtual address register has the same value as the internal ALU stage Program Counter (PC). When the exception occurs, a status/cause bit in the control register is set so that software can figure out what caused the exception. Tracepoints in a delay slot will set the BD CP0 Cause Register bit as expected; however, it is up to the software to jump past the delay slot correctly (by subtracting 4 and re-executing the branch).

Extended CPO Cache Configuration Register

The CPO Cache Configuration Register, as described in an earlier section, contains several software controllable Force Cache Miss features that allow logic analyzers to interface to the R36100.

Cause and EPC Register Writes

The R36100 adds a control bit, which if asserted, enables writes to the CPO Cause register cause field and the CPO EPC.

JTAG Scan Path

All the inputs and outputs are boundary scannable.

Features specific to debug/emulators

It is envisioned that operating system debug kernels always echo MTCO writes. Addresses are reserved in a scheme that frames the following registers:

- 32 General Purpose Registers.
- 32 CPO Registers (only 16 presently used for R3000 systems).
- 64 CP1 Floating Point Registers (only 32 presently used for R3000 systems).

Physical Address	Register
FFFF_8F68	K0 \$26
FFFF_8F6C	K1 \$27
FFFF_8F8C	CPO \$3 Config

Table 19.1 Reserved Emulator Addresses.

Pin Descriptions**Debug/emulator and Diagnostic Signals****DiagCache/UnCache** **Output**

Diagnostic Cached versus Uncached and Burst Miss Address 3: An output signal specifying cacheability type attribute of external system bus transactions. Signal is low during Uncached references, high during Cached ones. During the second clock of burst reads, outputs the miss address 3. The first and remainder clocks output cached versus uncached.

DiagInst/Data **Output**

Diagnostic Instruction versus Data Status and Burst Miss Address 2: An output signal specifying data type attribute of external system bus transactions. Signal is low during Instruction references, high during Data ones. During the second clock of burst reads, outputs the miss address 2. The first and remainder clocks output instruction versus data. Internal DMA transactions are always data transactions.

DiagRun **Output**

Diagnostic Run: A pseudo-synchronized active low output version of the internal CPU core RunN signal.

DiagBranchTaken **Output**

Diagnostic Branch Taken: A pseudo-synchronized active low output signal indicating when a branch is taken (same as the R3041A).

DiagJRorExe **Output**

Diagnostic Jump Register or Exception: A pseudo-synchronized active low output signal indicating when a jump register instruction is executed or an exception is taken.

DiagInternalWr **Output**

Diagnostic Internal Resource Write: An active low output signal indicating that an MTC0 instruction to register 3 was executed. This signal is used to indicate to the debug/emulator that it may want to interrogate the R36100 to find out if a control register that may have an effect on debug/emulator interpretation was altered.

DiagTriState **Input**

Diagnostic Tri-State: An input signal that when asserted low causes all outputs to tri-state. Can be used to:

1. Disable target board CPU during emulation.
2. Disable CPU during in-circuit manufacturing testing.

DiagInstCacheWrDis **Input**

Diagnostic Instruction Cache Write Disable: An active low input signal that disables instruction cache misses from updating the instruction cache. Meant to be asserted after $\overline{\text{DiagFICM}}$ and an instruction miss.

DiagFCM **Input**

Diagnostic Force Instruction and Data Cache Miss: An active low input signal causing all instructions and data loads (except internal partial word store reads) to miss the cache and do an external system bus read. In this mode no newly initiated read cache misses are written into the cache. During the assertion of $\overline{\text{DiagFCM}}$, internal generation of 'AckN' is disabled.

Note: Although emulators typically assert this pin during functional operation, the non-emulator user should either assert or not assert this pin during power-up, and leave it continuously asserted or not asserted.

DiagIntDis **Input**

Diagnostic Interrupt Disable: An active low input signal when asserted, causes all external and internal interrupts to be disabled.

DiagNoCS **Output**

Diagnostic No Chip Select: An active low input signal asserted concurrent with $\overline{\text{SysALEn}}$ indicating that no external or internal chip select was activated for this read or write.

DiagInternalDmaBusGnt **Output**

Diagnostic Internal DMA Channel Bus Grant: An active low output signal asserted whenever one of the four internal DMA channels receives a bus grant. This signal can be gated with a peripheral chip select to distinguish between a peripheral control register access versus a DMA access.

JTAG Signals

JtagClkIn (TCLK)	Input
JtagModeSelect (TMODE)	Input
JtagDataIn (TDI)	Input
JtagDataOut (TDO)	Output
JtagReset (TRES*)	Input

See IEEE Specification 1149, "Standard Test Access Port and Boundary Scan Architecture" for more information on the specification of these signals. At a minimum, if JTAG is unused in the system application, TMODE must be pulled/tied high for at least 5 TCK rising edges and/or TRES* must be pulled/tied low in order to ensure that the JTAG circuitry is properly reset into an inactive state. TRES* must be pulled/tied low and TMODE must be pulled/tied high at all times after SysReset is asserted.

The R36100 supports the following JTAG Instructions: EXTEST; BYPASS; and SAMPLE/PRELOAD.

Register Descriptions

Note that Big Endian software must offset these addresses by b'10 (0x2).

Phys. Address	Description
0xFFFF_E500	Tracepoint Control Register
0xFFFF_E504	Debug Control Register
0xFFFF_E510	TraceLSB(0) Address Register
0xFFFF_E514	TraceMSB(0) Address Register
0xFFFF_E520	TraceLSB(1) Address Register
0xFFFF_E524	TraceMSB(1) Address Register

Table 19.2 Debug Interface Register Address Assignments

MSB Debug Tracepoint Address Register LSB Debug Tracepoint Address Register (‘DebugTraceAddrReg’)

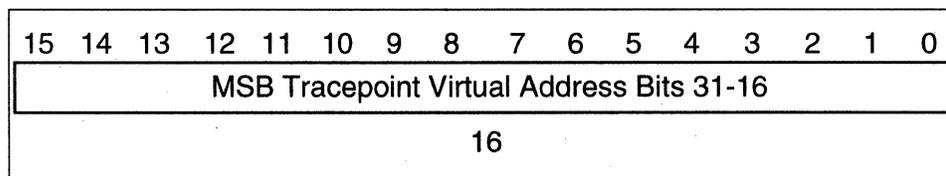


Figure 19.1 Debug Tracepoint Address Register (‘MSB DebugTraceAddrReg’).

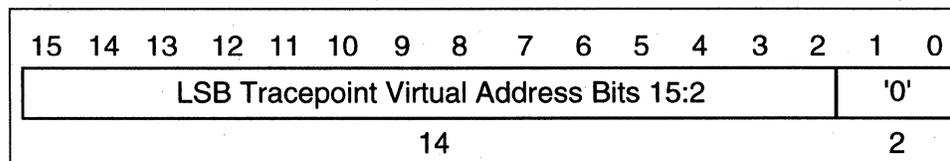


Figure 19.2 Debug Tracepoint Address Register (‘LSB DebugTraceAddrReg’).

When the Debug Tracepoint Virtual Address Register MSB and LSB matches the internal program counter (ALU pipeline stage), and the feature is enabled via the debug tracepoint control register, an exception is taken and a debug tracepoint control register cause bit is set.

Bit	Description
15:2	Tracepoint Virtual Address
1:0	Reserved '0'

Table 19.3 Debug Tracepoint Address Register ('DebugTraceAddrReg') Bit Assignments.

Debug Tracepoint Control Register ('DebugTraceControlReg')

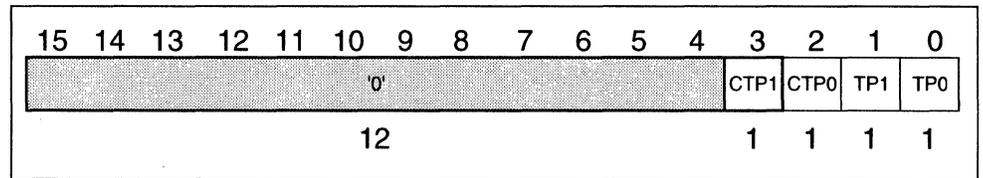


Figure 19.3 Debug Tracepoint Control Register ('DebugTraceControlReg').

The Debug Tracepoint Control Register is used to access and control tracepoint and single step functions.

Bit	Description
15:4	Reserved '0'
3	CTP1
2	CTP0
1	TP1
0	TP0

Table 19.4 Table Debug Tracepoint Control Register ('DebugTraceControlReg') Bit Assignments.

Reserved Low ('0') Field:

Must be written to '0' for future compatibility. Value when read is undefined.

Cause is Tracepoint ('CTP') Field:

After getting an exception, if the CTP field is found to be an active '1', the exception handler should acknowledge the exception by writing a '0' to the CTP bit. There are two fields: branch taken/branch not taken'.

Value	Action
'1'	Cause of exception is Tracepoint.
'0'	Cause of exception is not Tracepoint (default).

Table 19.5 Cause is a Tracepoint ('CTP') Field Encoding.

Tracepoint ('TP') Field:

Value	Action
'1'	Tracepoint On. Forces CPU to allow tracepoint register to activate if the Program Counter matches the Tracepoint Address Register.
'0'	Tracepoint Off (default).

Table 19.6 Tracepoint Enable ('TP') Field Encoding.

Debug Control Register ('DebugControlReg')

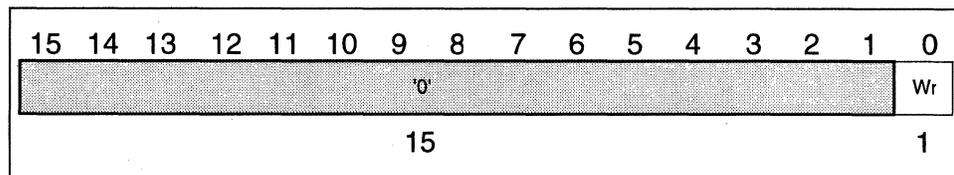


Figure 19.4 Debug Control Register ('DebugControlReg').

Bit	Description
15:1	Reserved '0'
0	Wr

Table 19.7 Debug Control Register ('DebugControlReg') Bit Assignments.

Reserved Low ('0') Field:

Must be written to '0' for future compatibility. Value when read is undefined.

Writability ('Wr') Field:

Value	Action
'1'	Allow CPO Cause Bits and EPC Register to be written.
'0'	CPO Cause Bits and EPC Register are read only.

Table 19.8 Writability ('Wr') Field Encoding.

Initializing SysClk for Test

Another feature for board level testing is the ability to initialize the phase of $\overline{\text{SysClk}}$ to its high phase. A low to high transition on $\overline{\text{Reset}}$ will cause the internally synchronized (delay of less than or equal to 2 clocks) version of $\overline{\text{Reset}}$ to always set $\overline{\text{SysClk}}$ high during its next phase. Thus the state of $\overline{\text{SysClk}}$ can be deterministically controlled within a known number of ClkIn transitions. The two cases are shown in Figure 19.5 and in Figure 19.6.

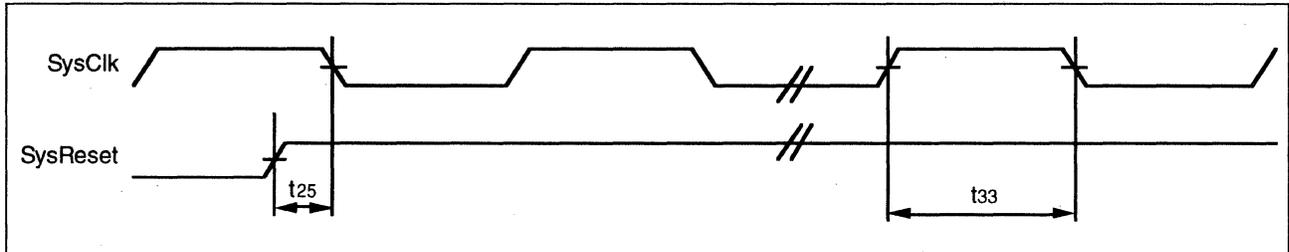


Figure 19.5 R36100 SysClk Phase Initialization Case A

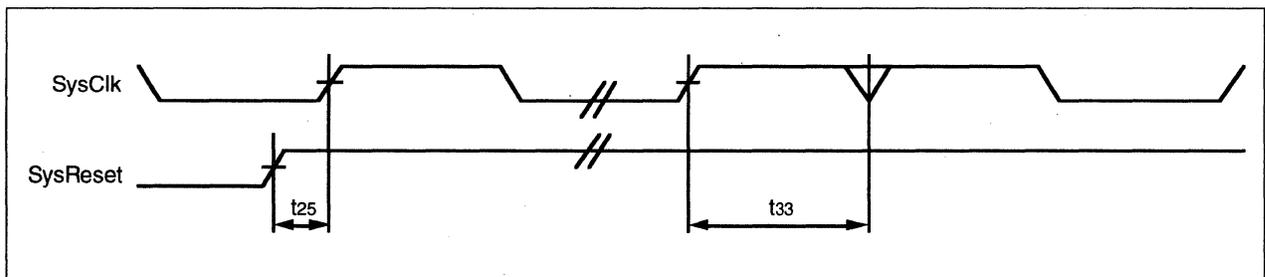


Figure 19.6 R36100 SysClk Phase Initialization Case B

Using Diag for Instruction Disassembly

The R36100 provides a Diagnosis pin which during its data phase outputs whether a read transaction is the result of an instruction fetch or the result of a data fetch. This information is independent of the information given during the address phase of whether or not the read was a result of a cached or uncached read. Note that this pin is undefined on writes; however, by necessity all writes must be data writes.



Integrated Device Technology, Inc.

2975 Stender Way
P.O. Box 58015
Santa Clara, CA 95052-8015
800-345-7015
FAX 408-492-8674

ELECTRONIC ACCESS

Internet: www.idt.com

E-Mail: info@idt.com

FAX-On-Demand: 800-9-IDT-FAX (in U.S.)
408-492-8391 (outside U.S.)



www.idt.com