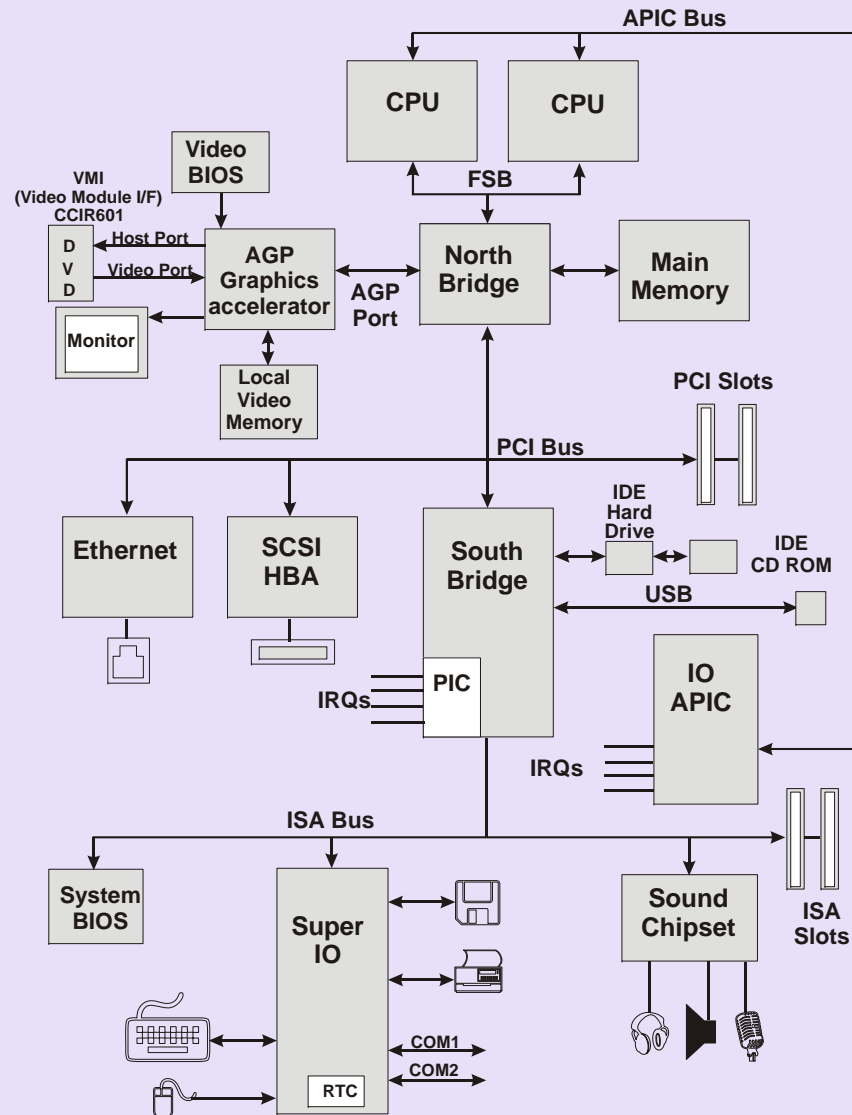




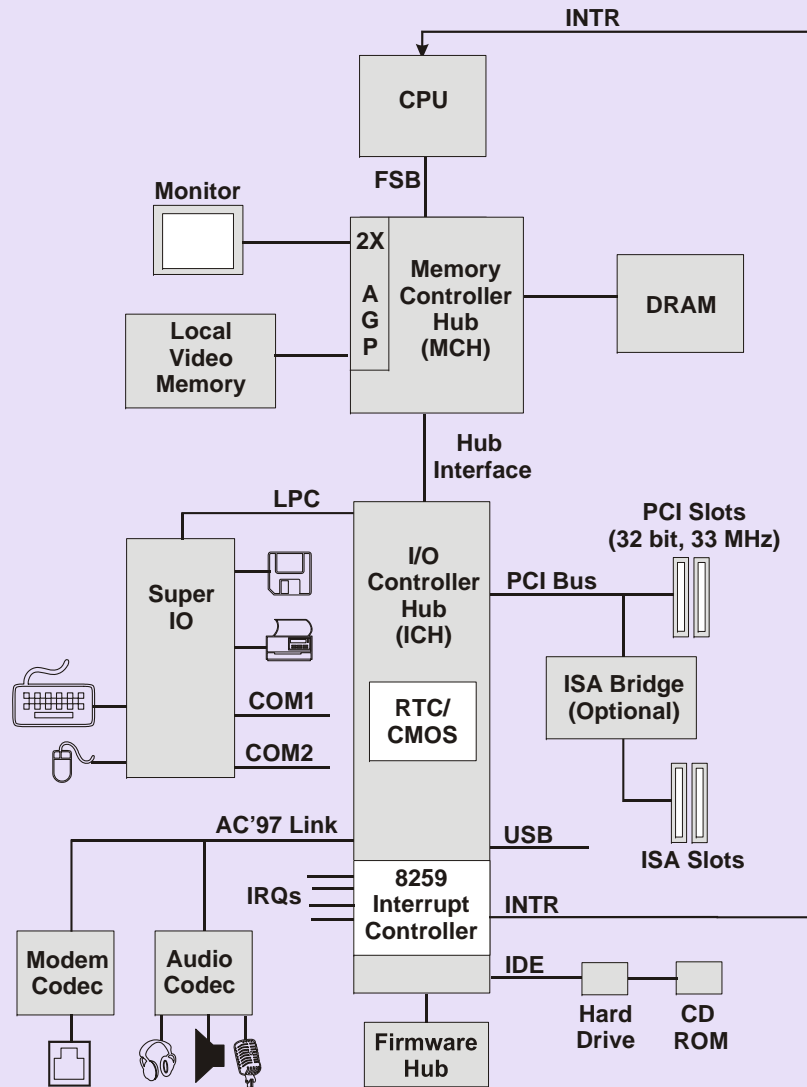
Conventional PCI Overview

Ravi Budruk
Senior Staff Engineer and Partner
MindShare, Inc.

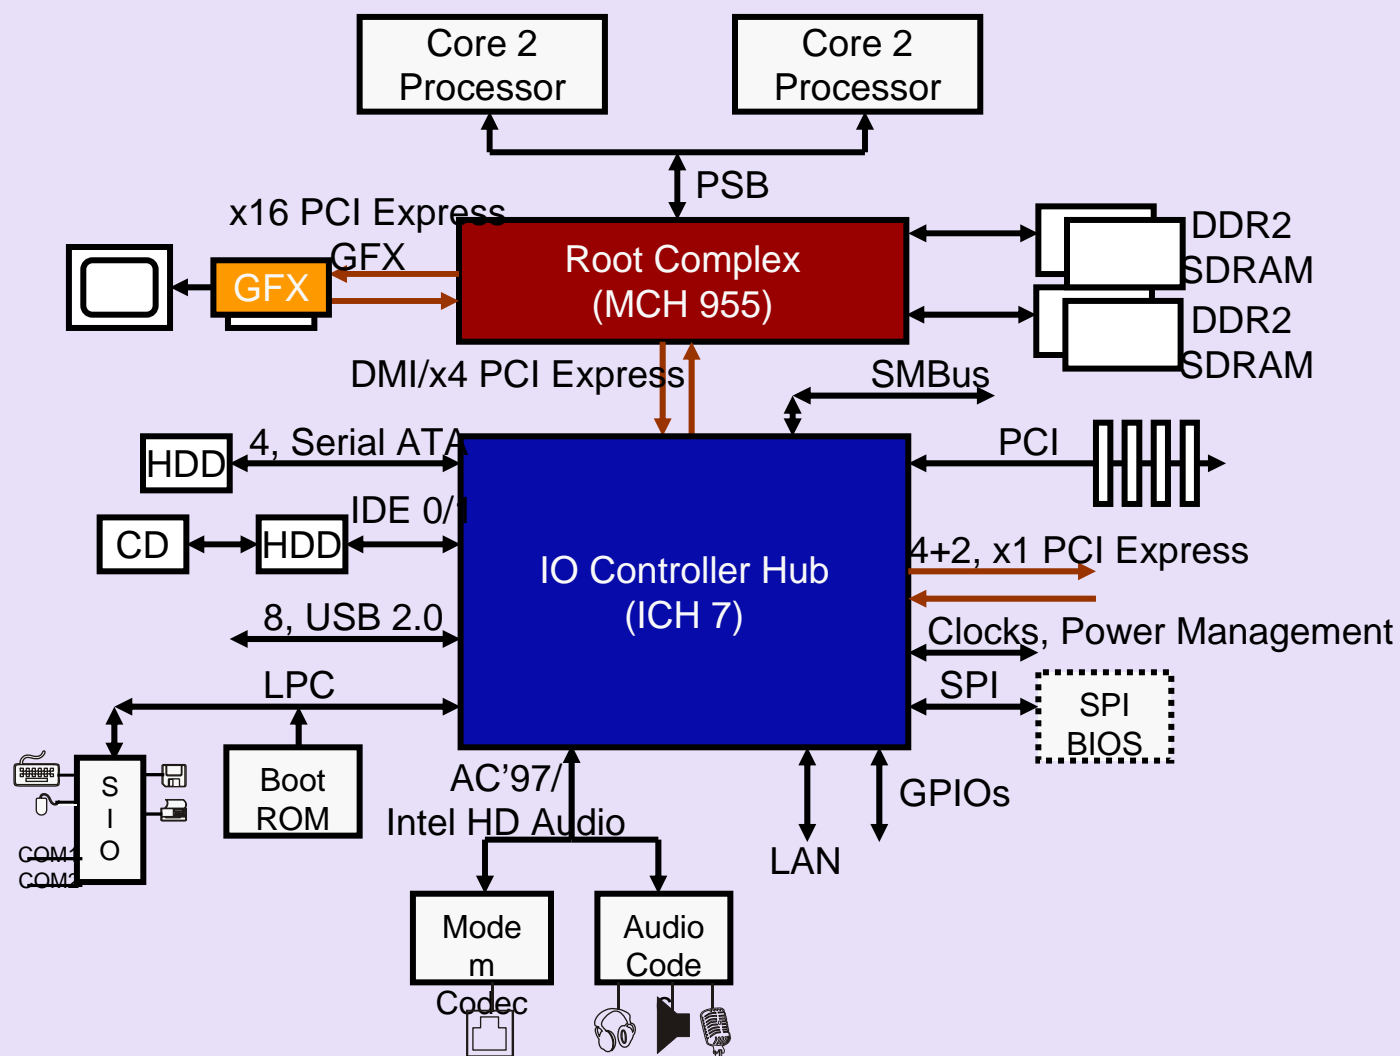
Example of Intel 440 Chipset Based PC System



Example of Intel 820 Chipset Based PC System



Example of Intel 955 Chipset Based PC System



PCI Bus Features


- Processor Independence
- Support for 256 functional devices per bus
- Support for up to 256 PCI buses
- Low power consumption
- Burst transactions supported
- 33MHz or 66MHz top frequencies
- 32- or 64-bit data bus width
- Access time as little as 2 clocks for writes, 3 clocks for reads

PCI Bus Features, continued


- Concurrent bus operation
- Bus master support
- Hidden bus arbitration
- Low pin count
- Transaction integrity checking
- Three address spaces:
 - ✓ Memory (4GB optionally 16EB),
 - ✓ I/O (64KB optionally 4GB), &
 - ✓ Configuration
- Auto configuration
- Software transparency
- Add-in cards with different dimensions

PC Address Space Mapping


Memory Map



IO Map



PCI Configuration Space




CFCh-CFFh

CF8h-CFBh

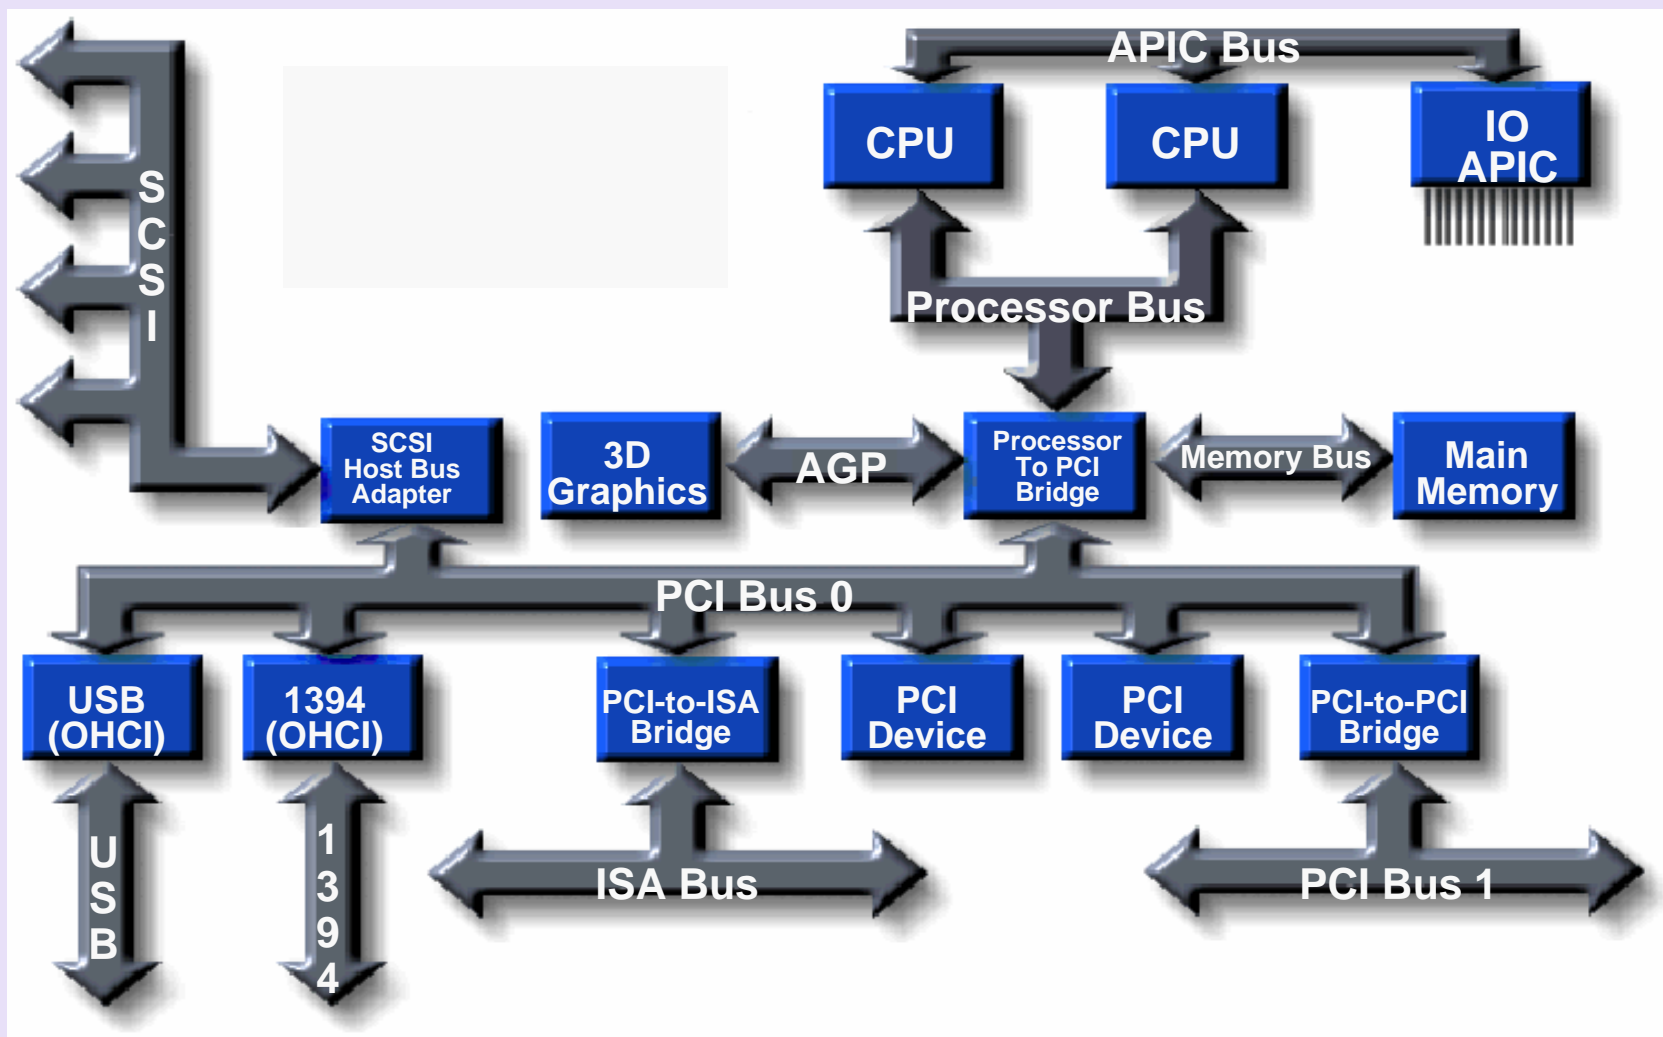
Intro to Burst Transfers & Protocol



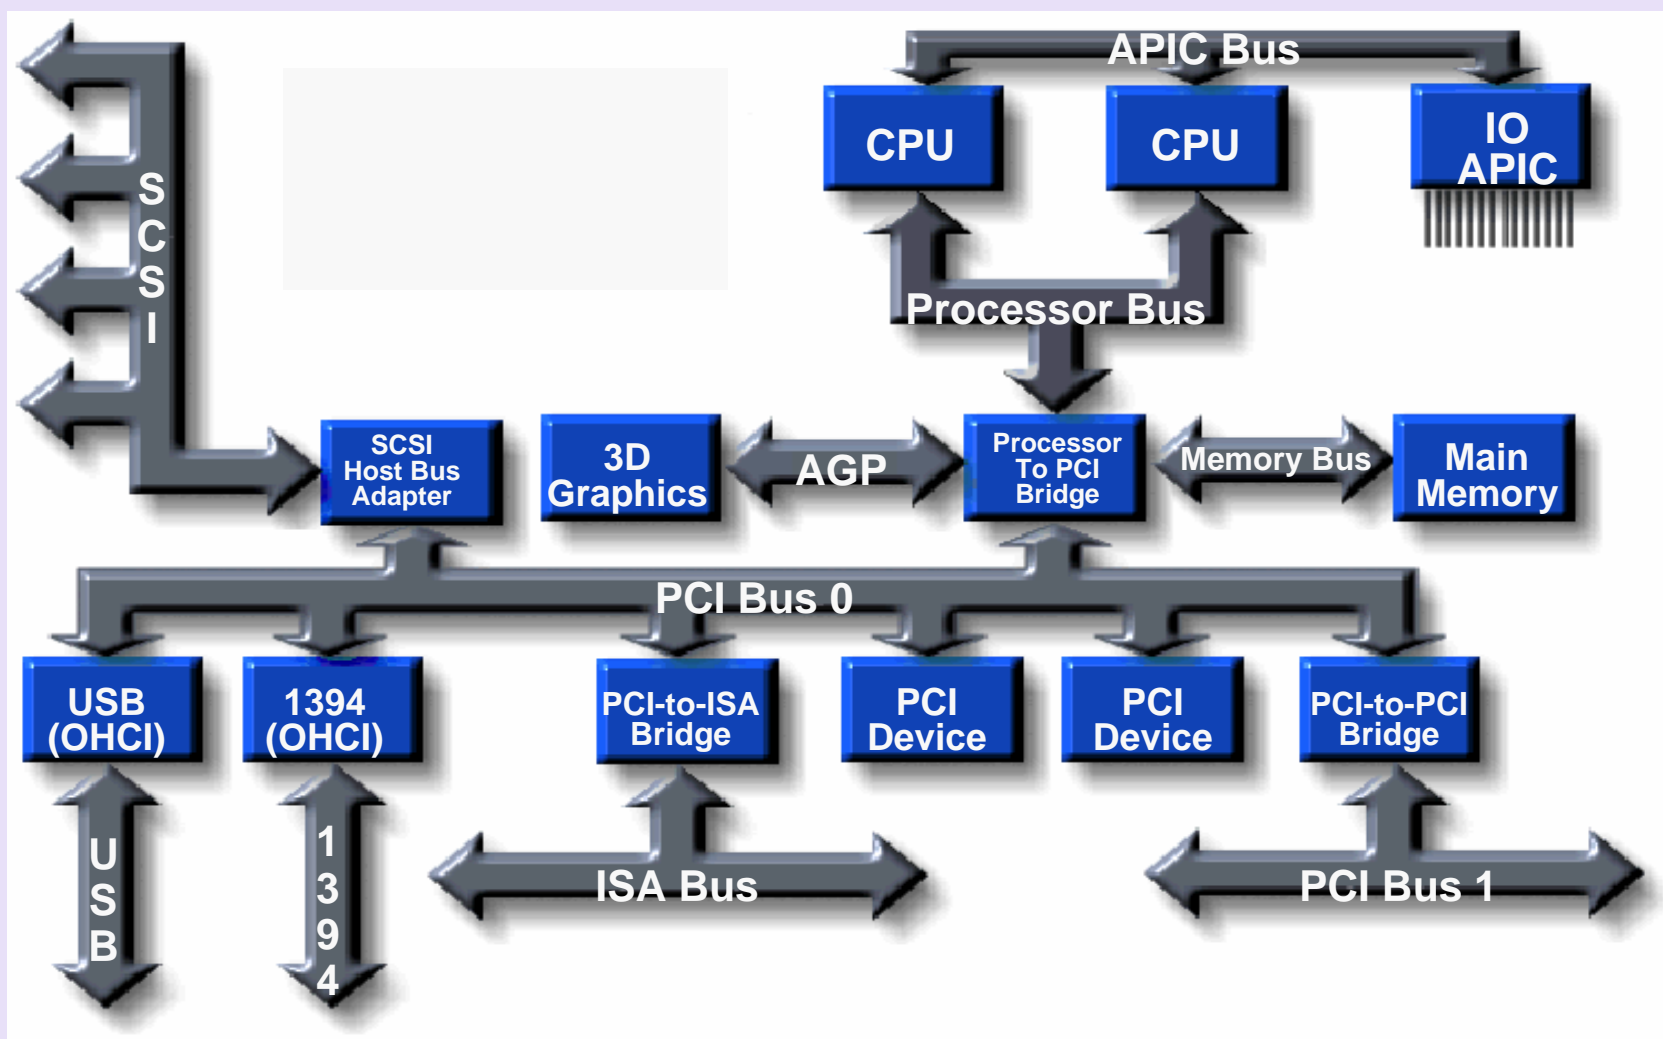
Burst Transfer




Initiator, Target, and Agents



Single vs Multi-Function Device




Clock




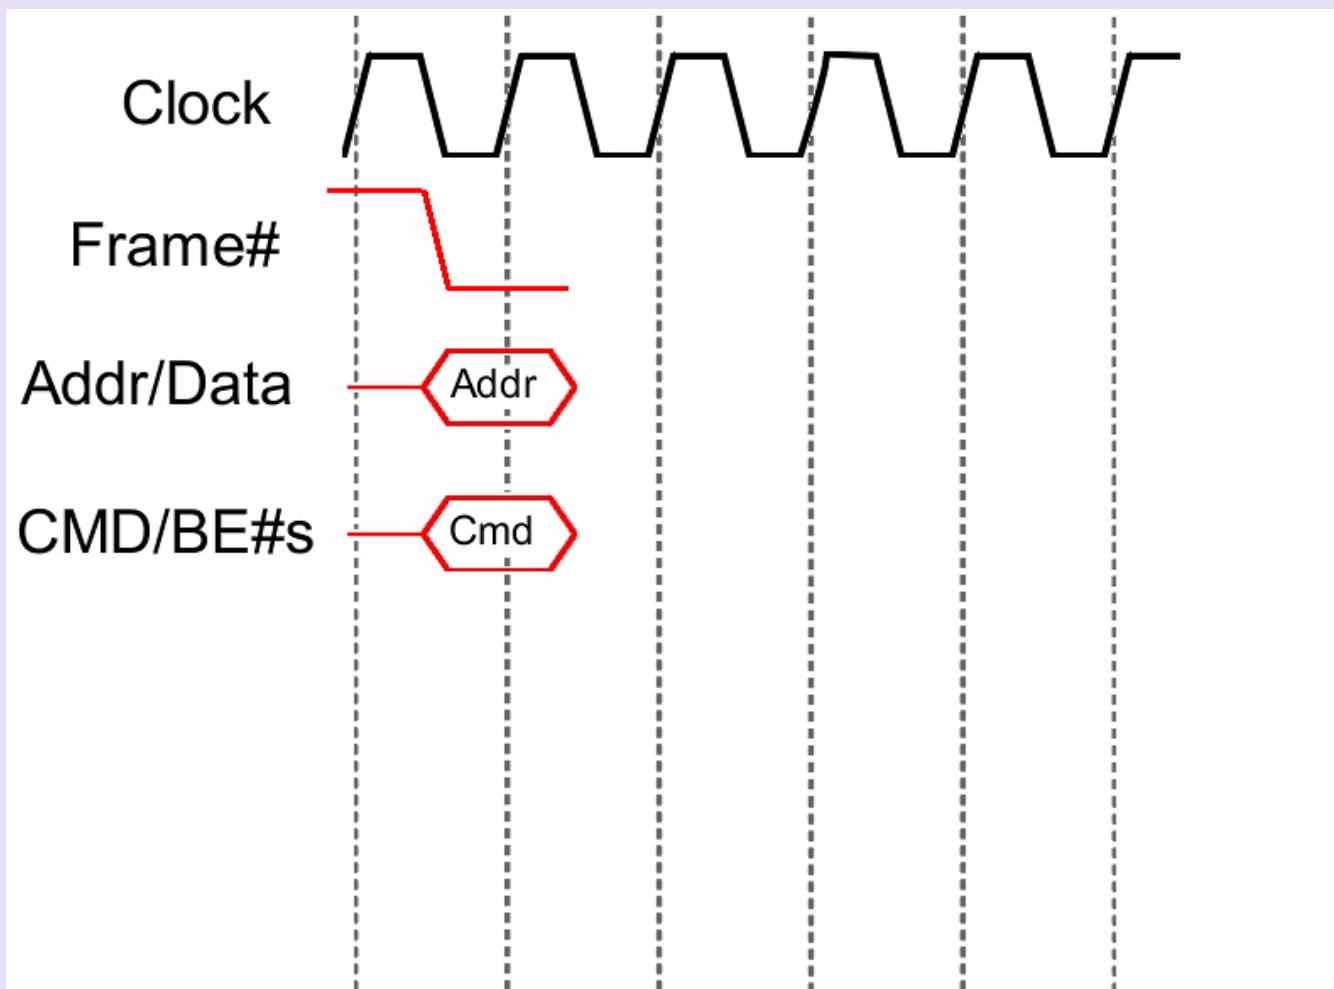
PCI Transaction Protocol Fundamentals


Clock

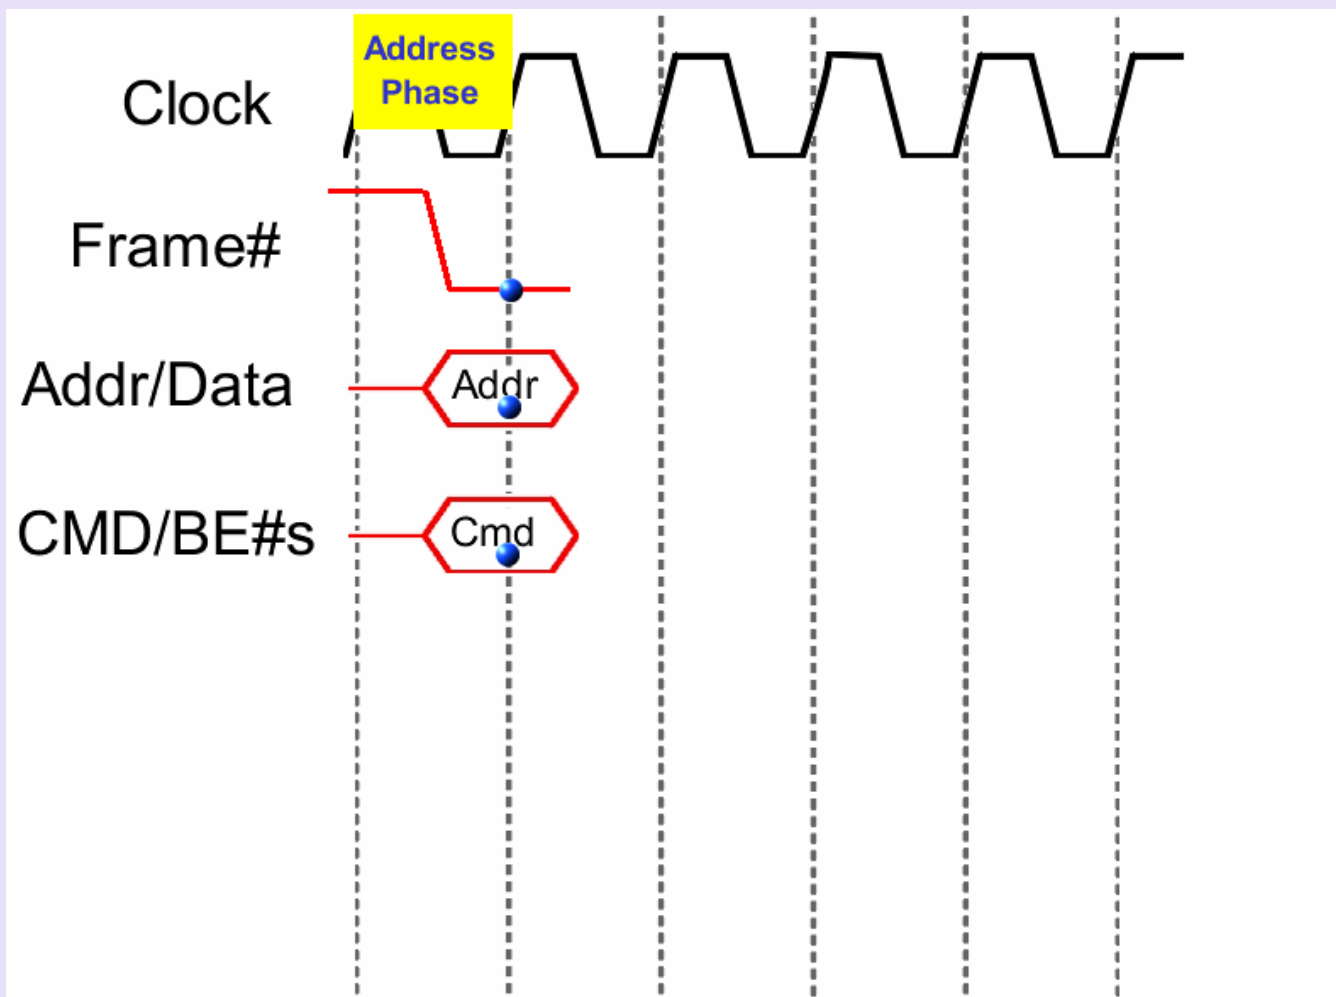


Address Phase

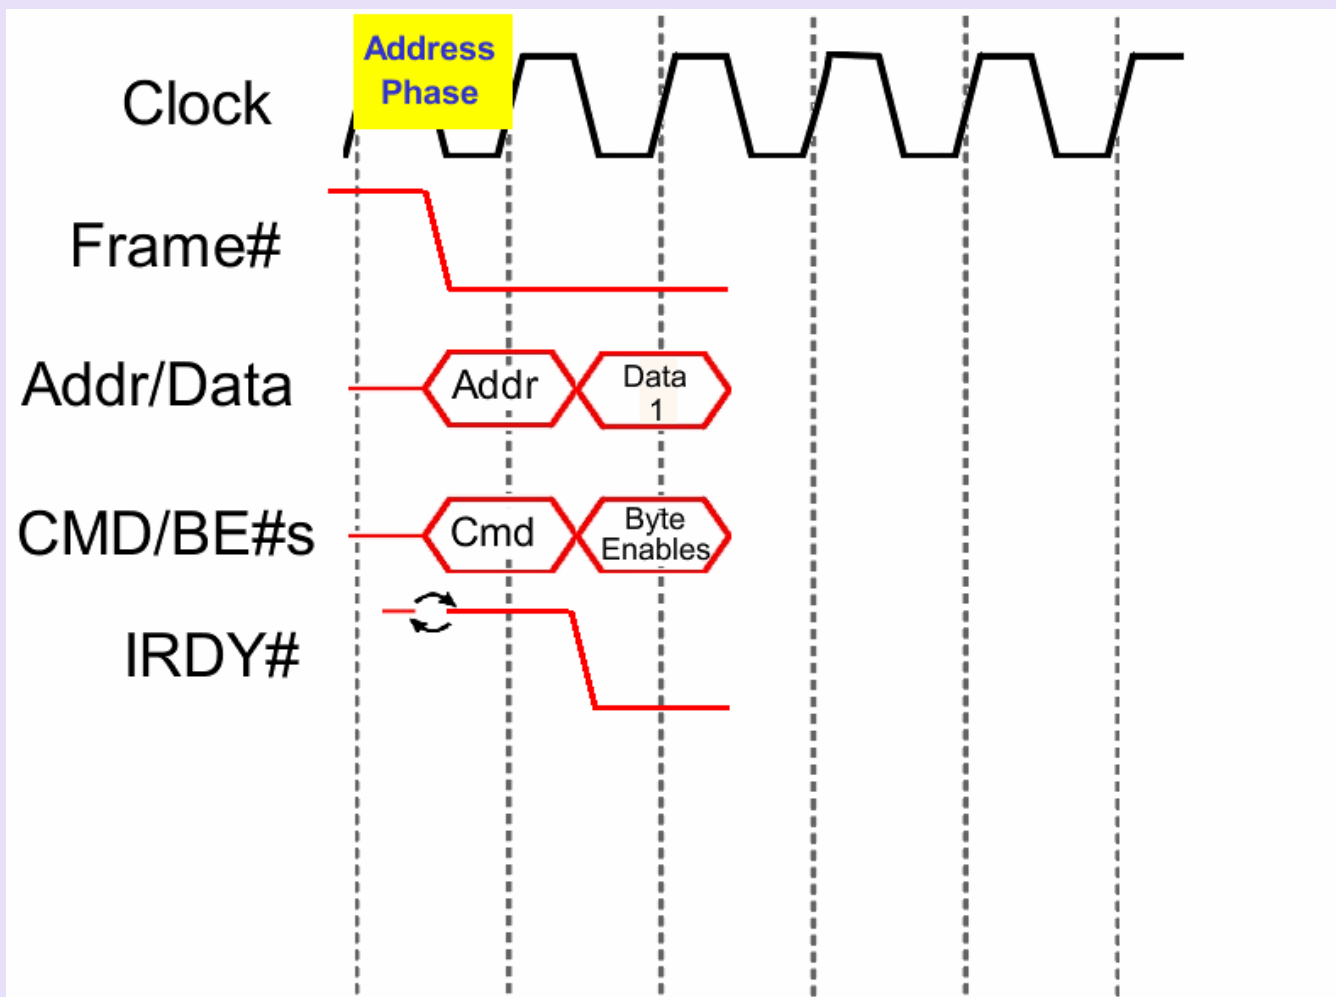




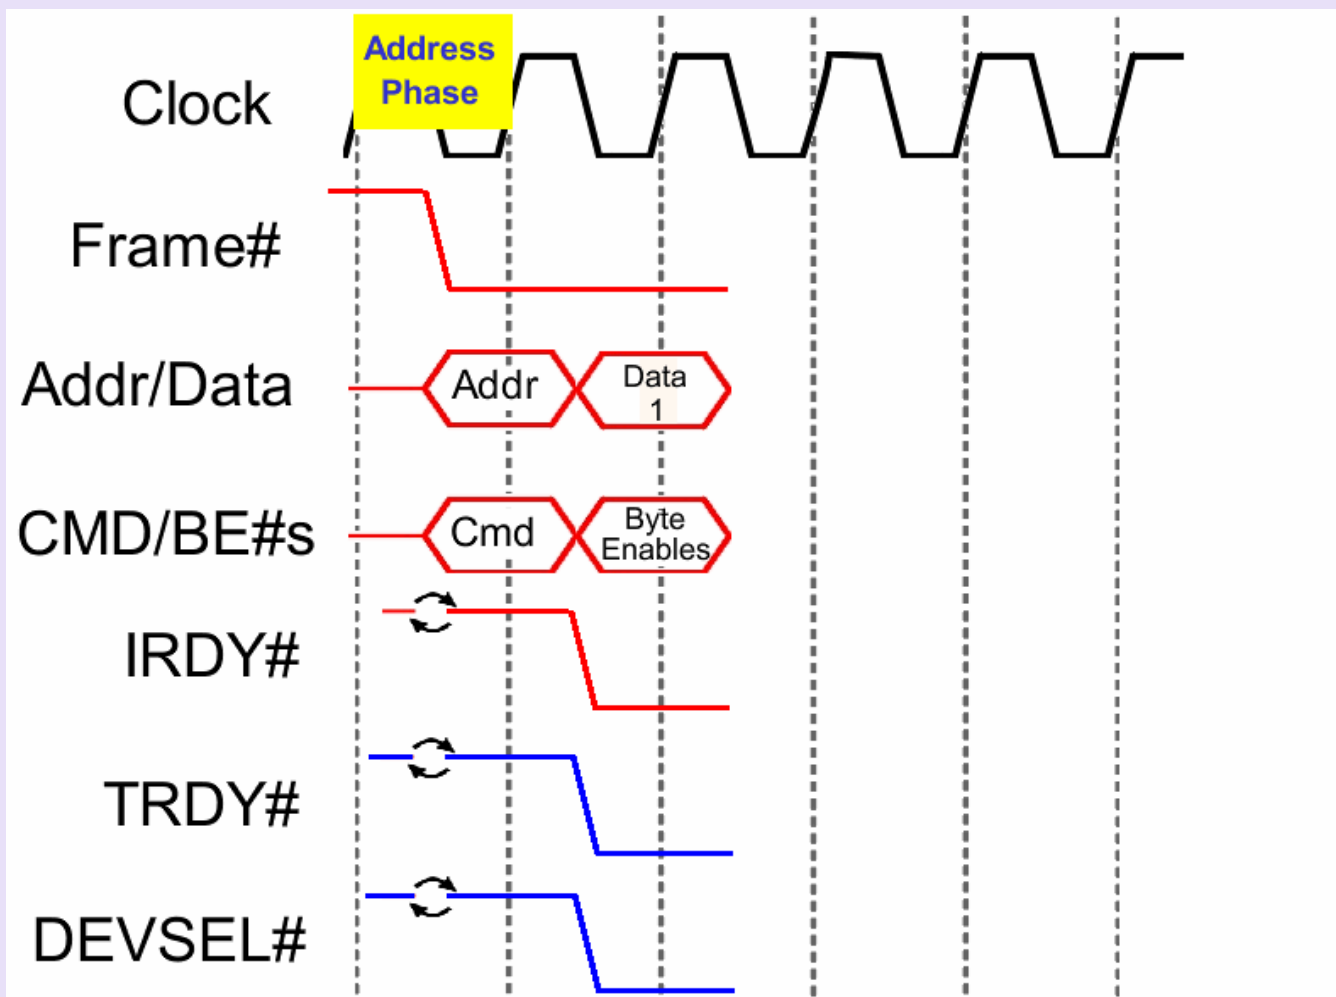





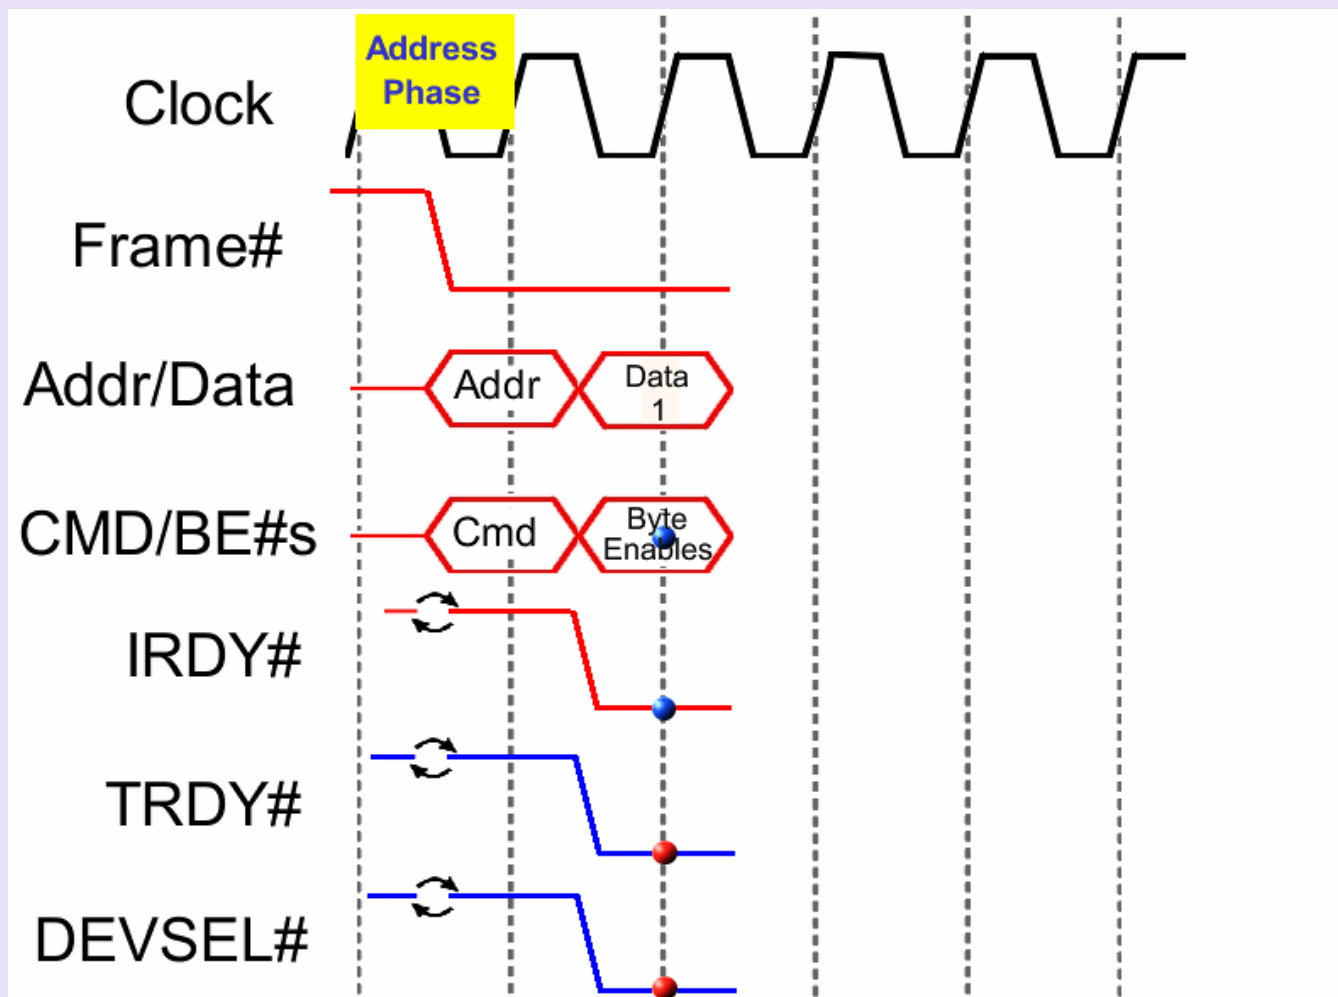
Data Phases

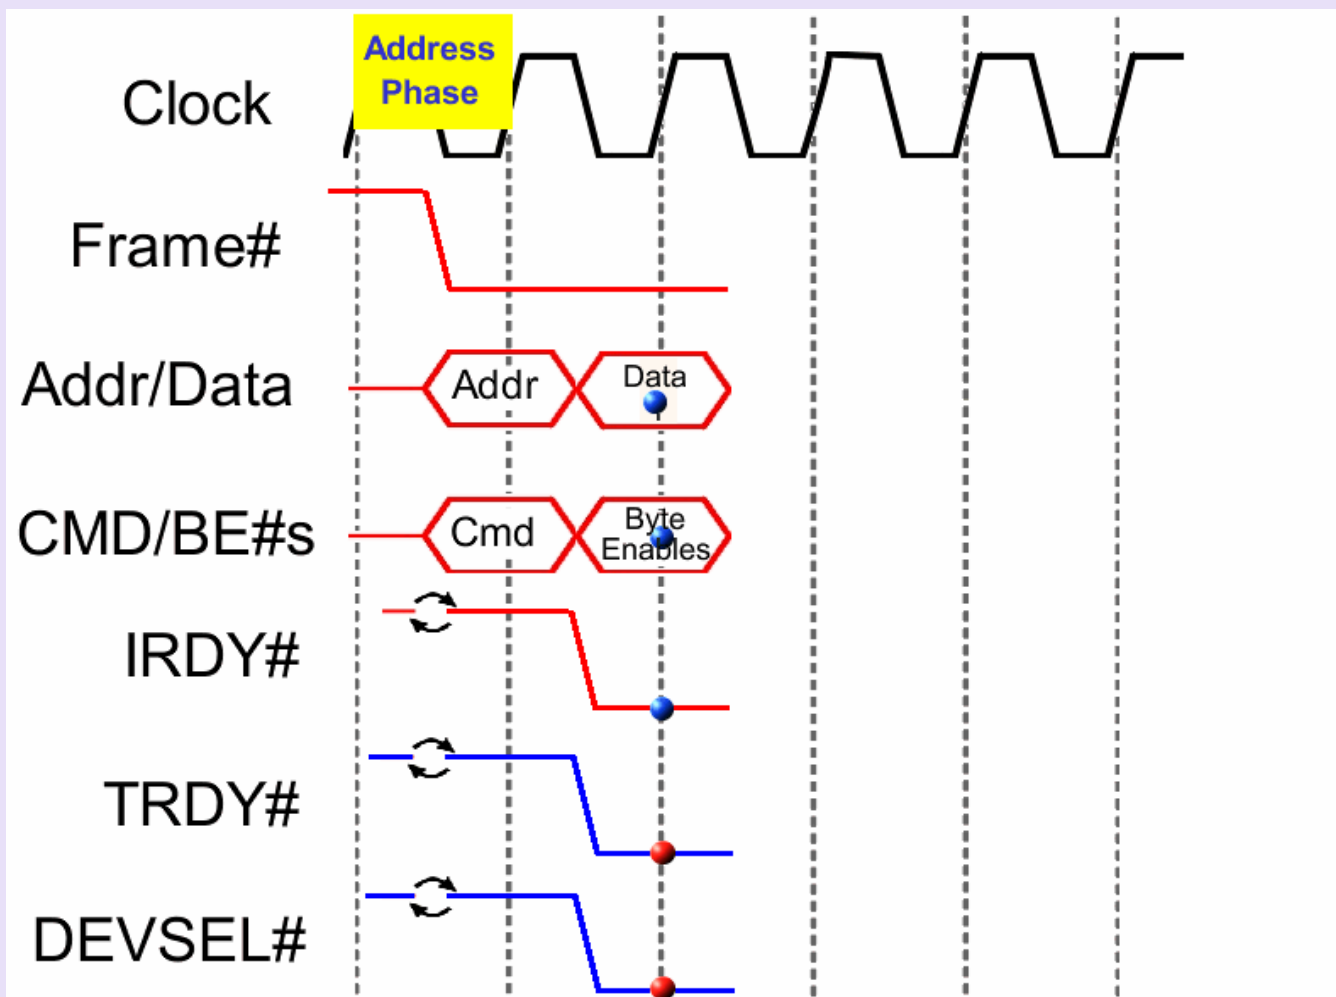


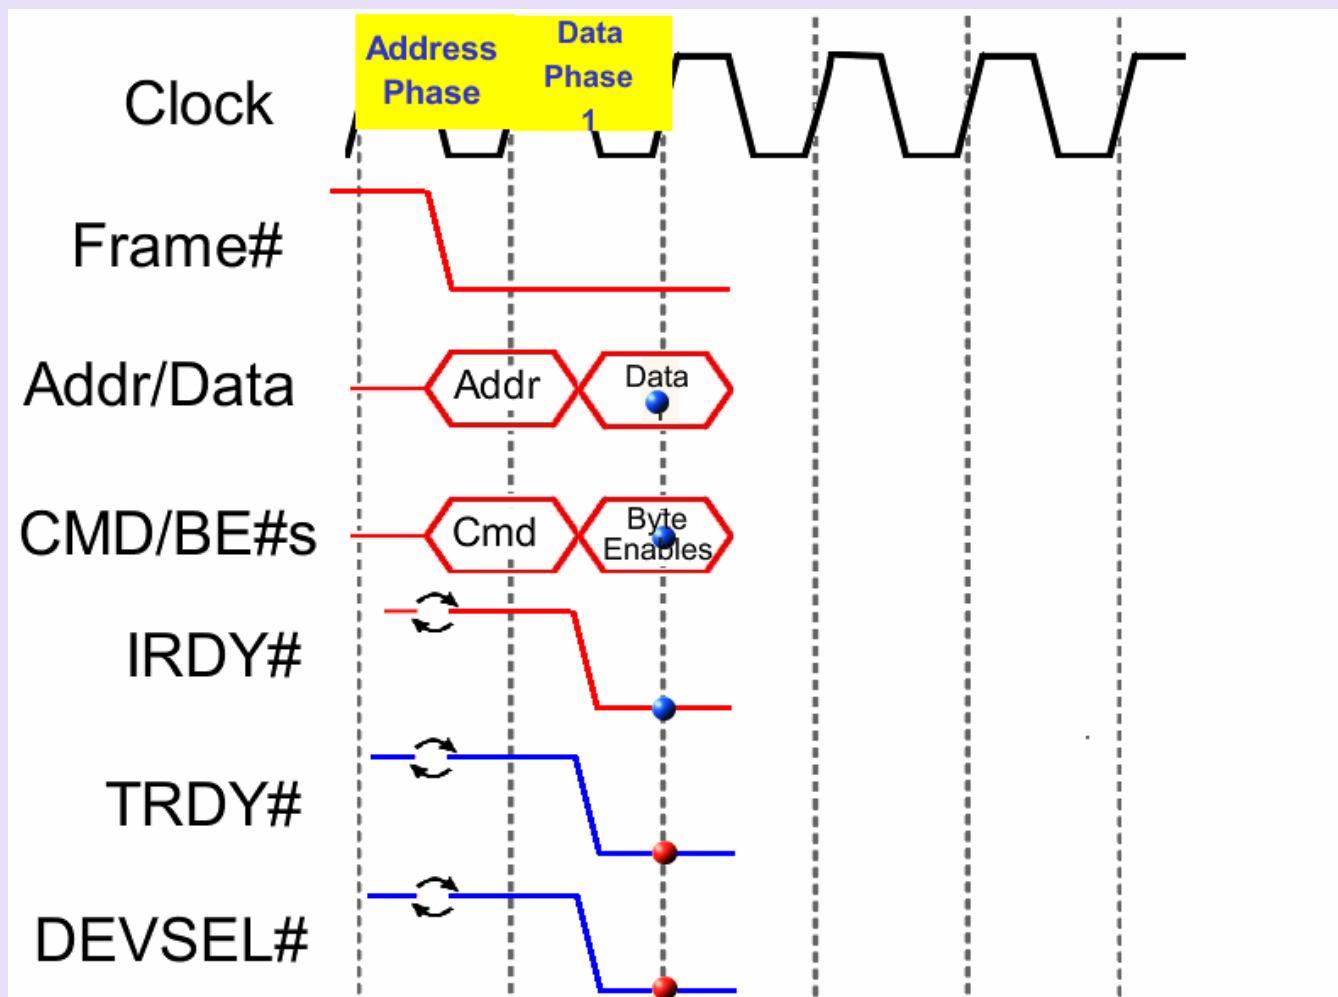
Claiming the Transaction

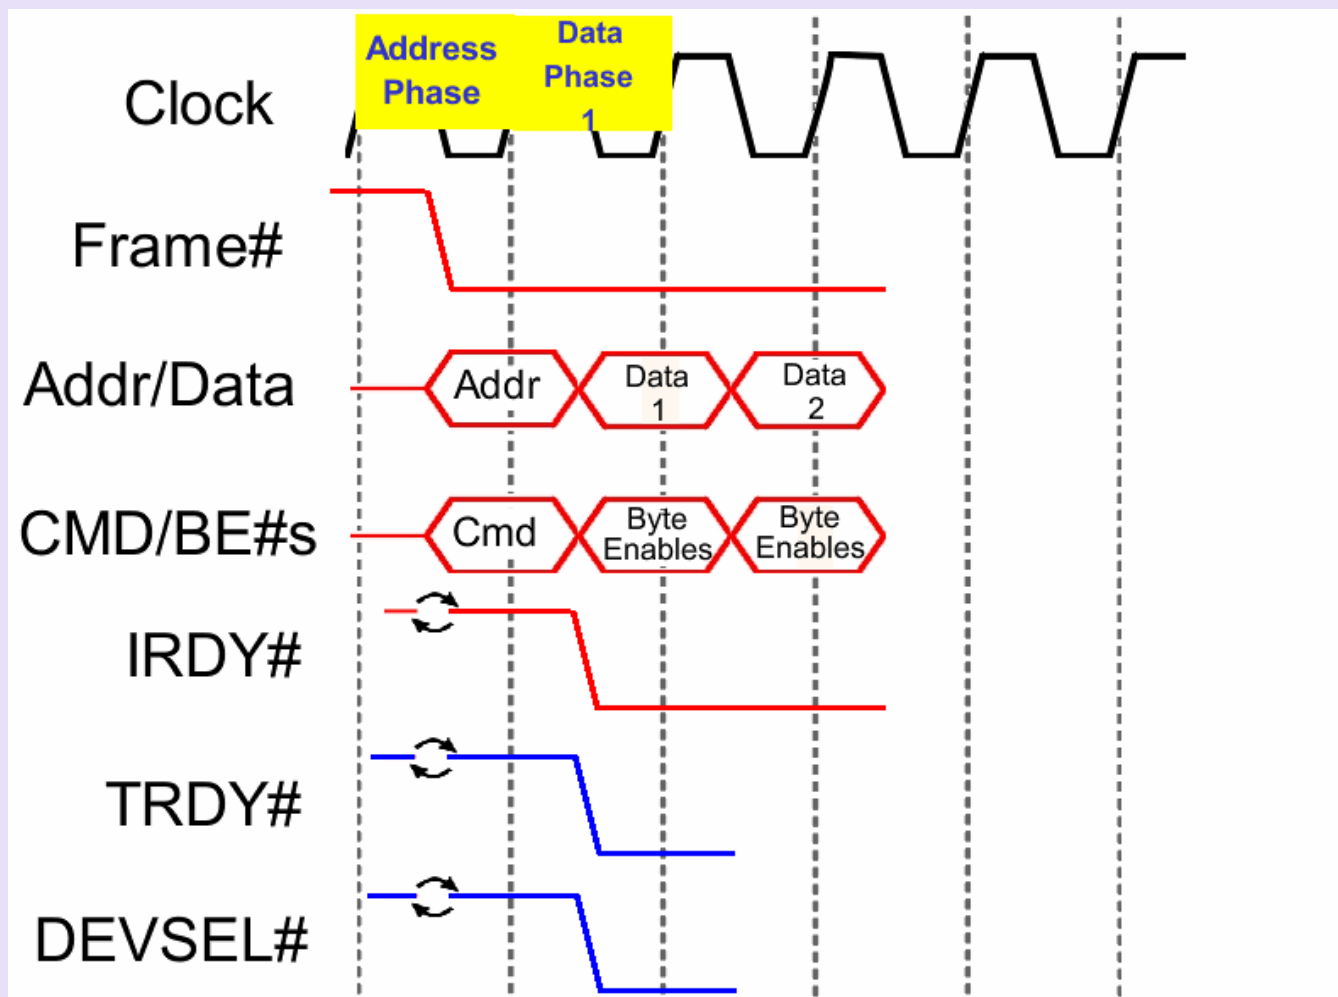


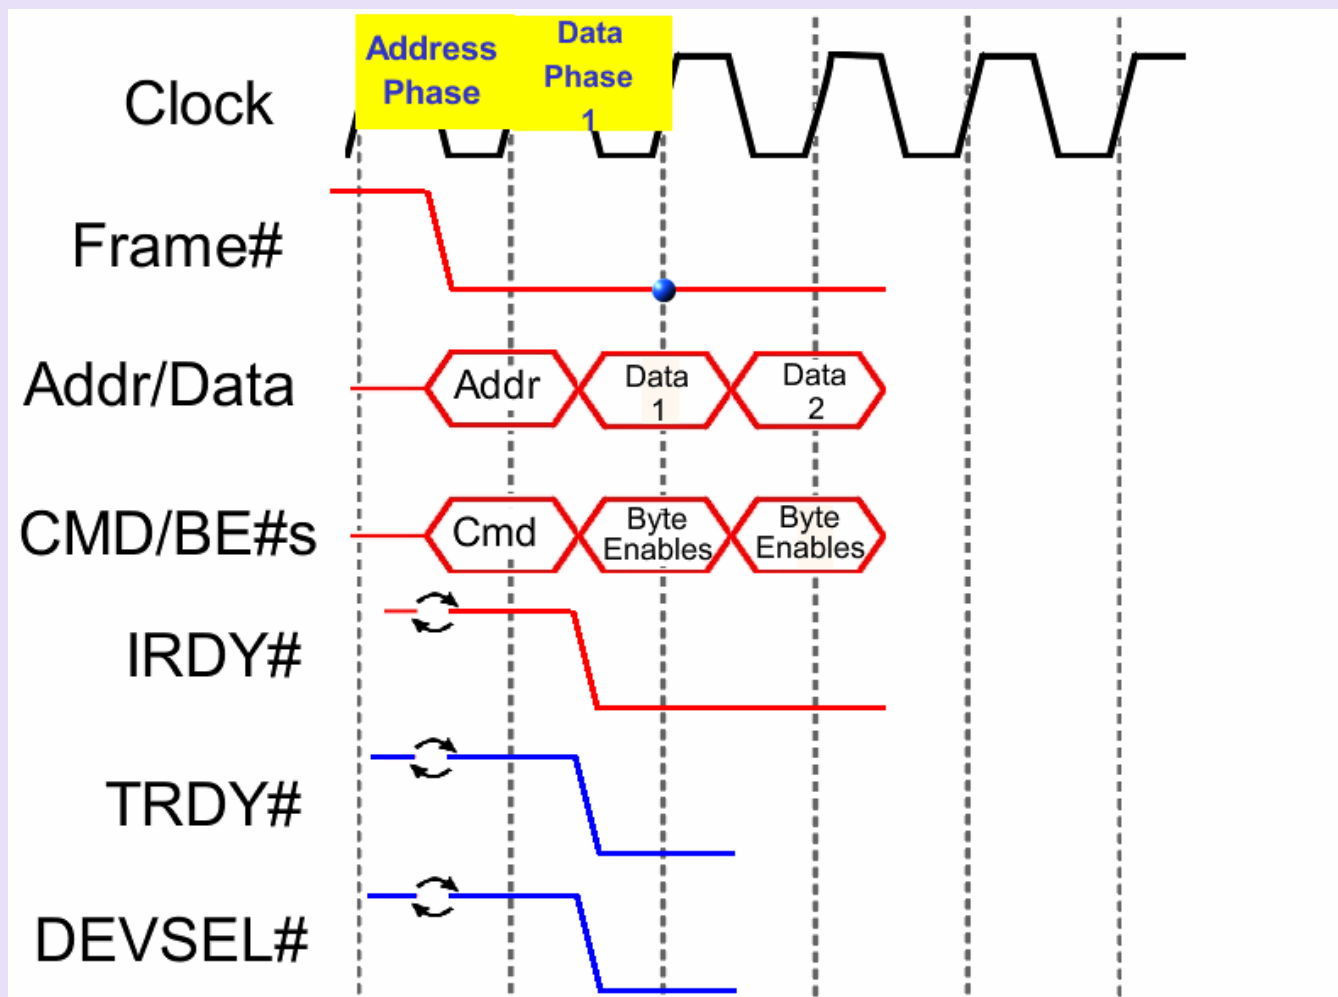


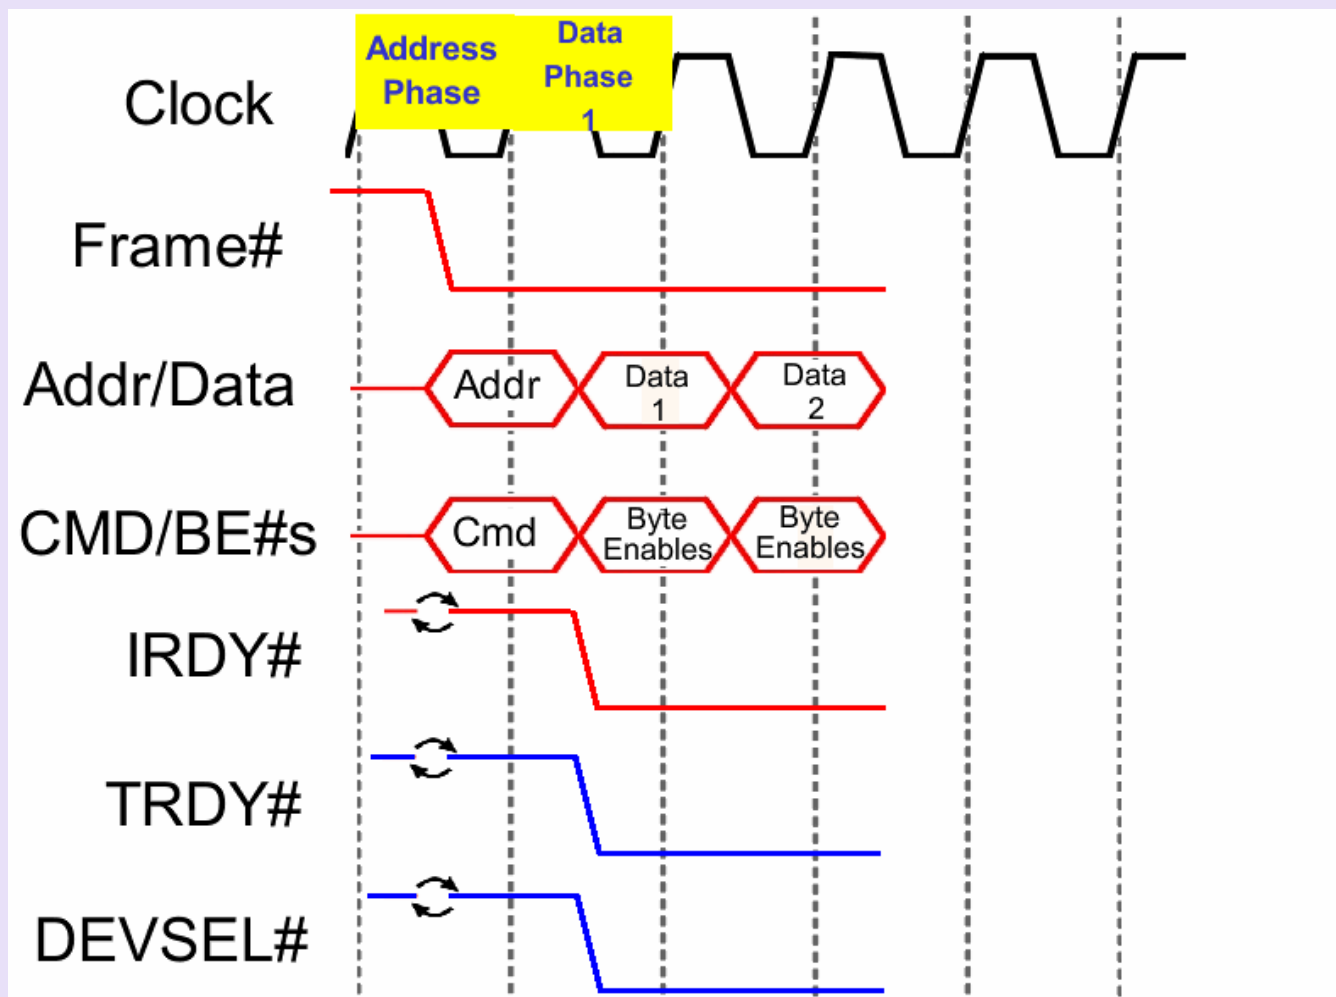








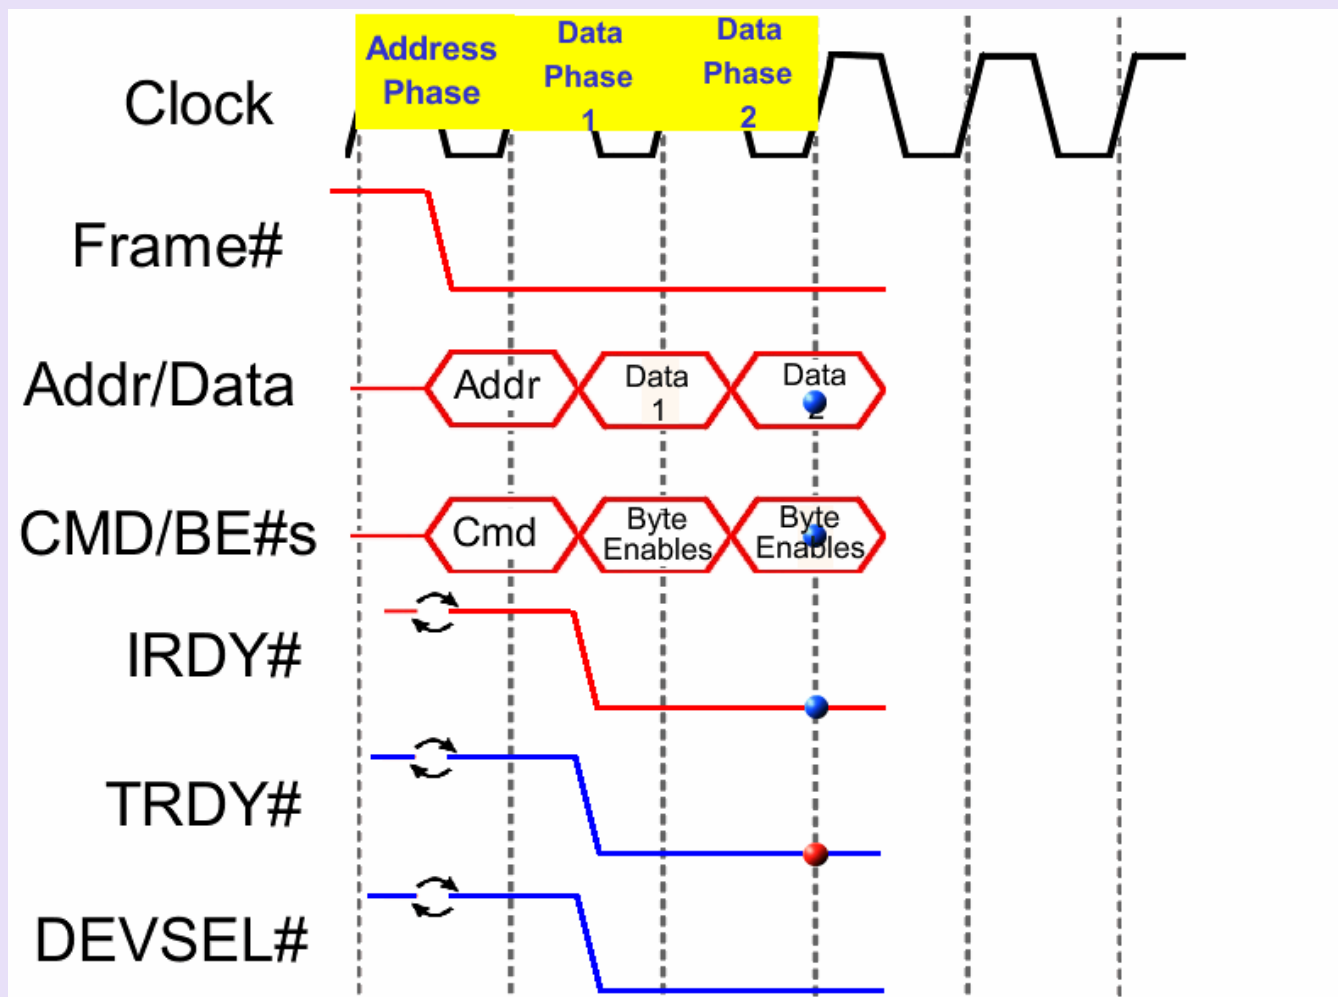







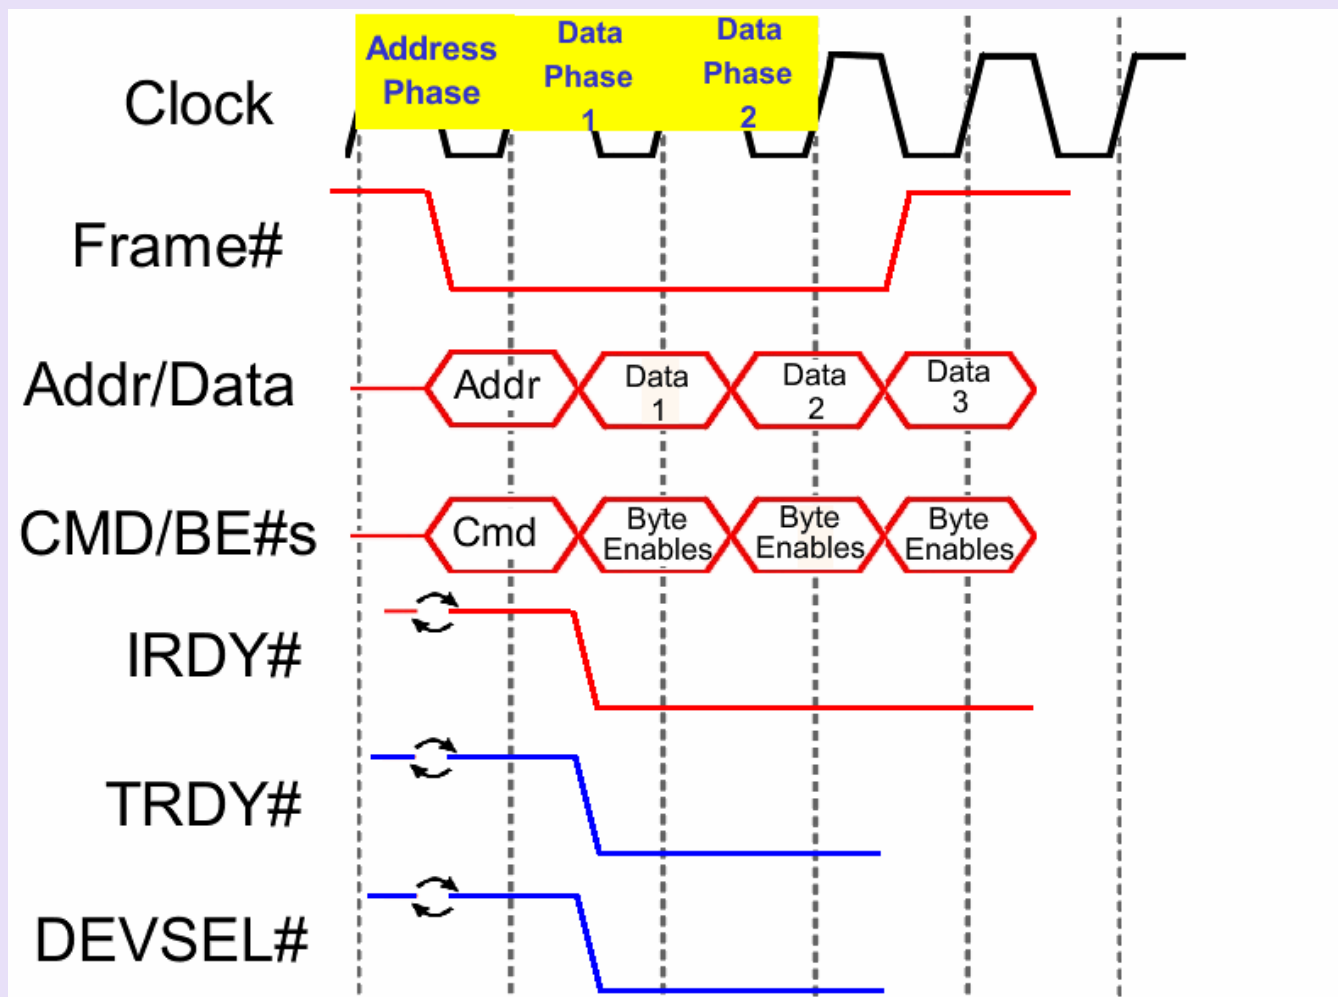


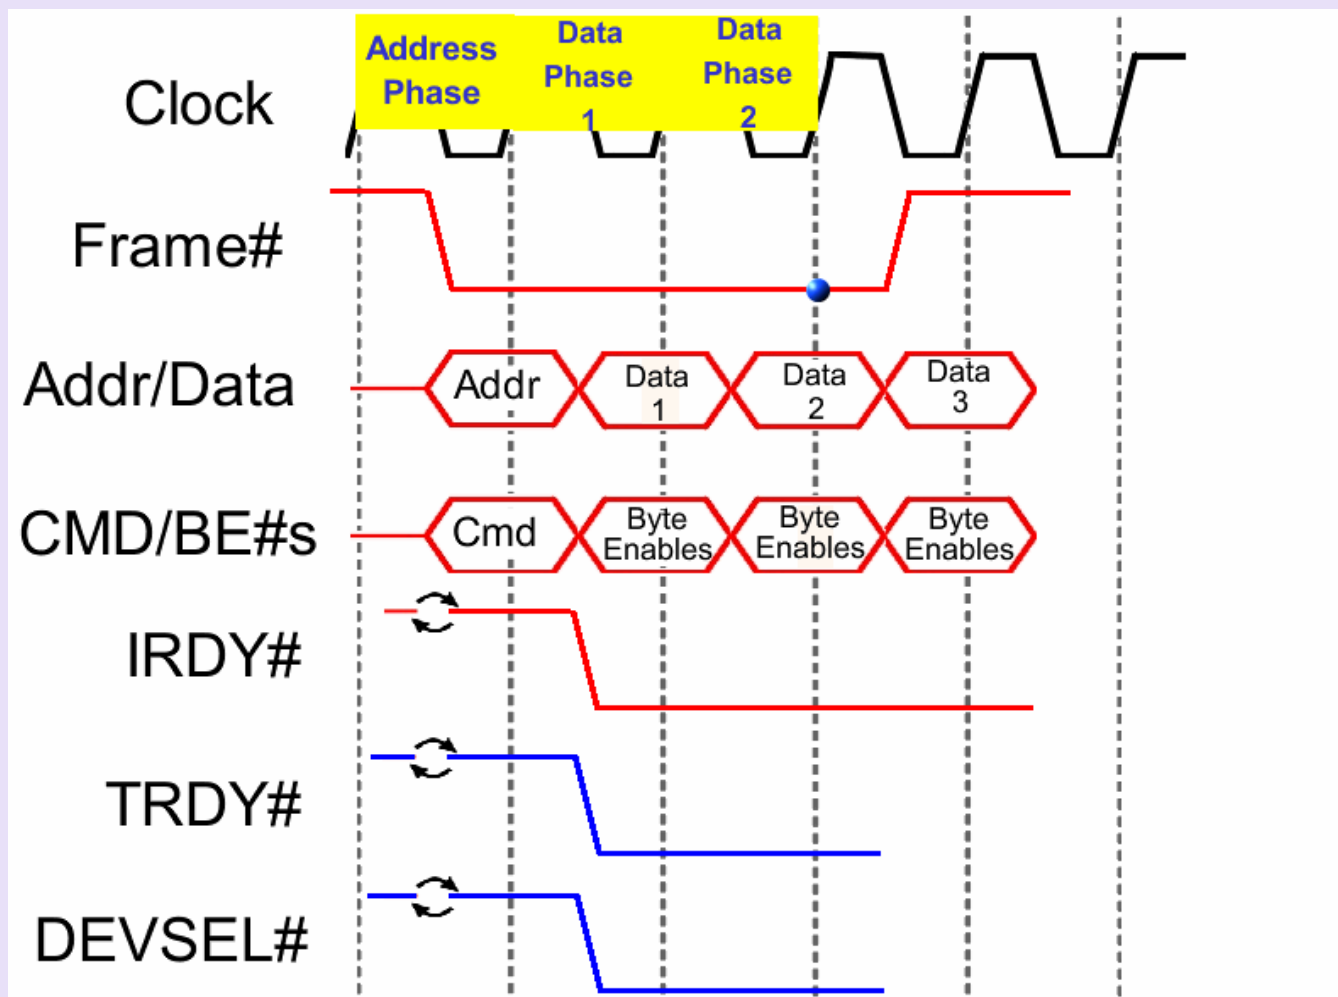


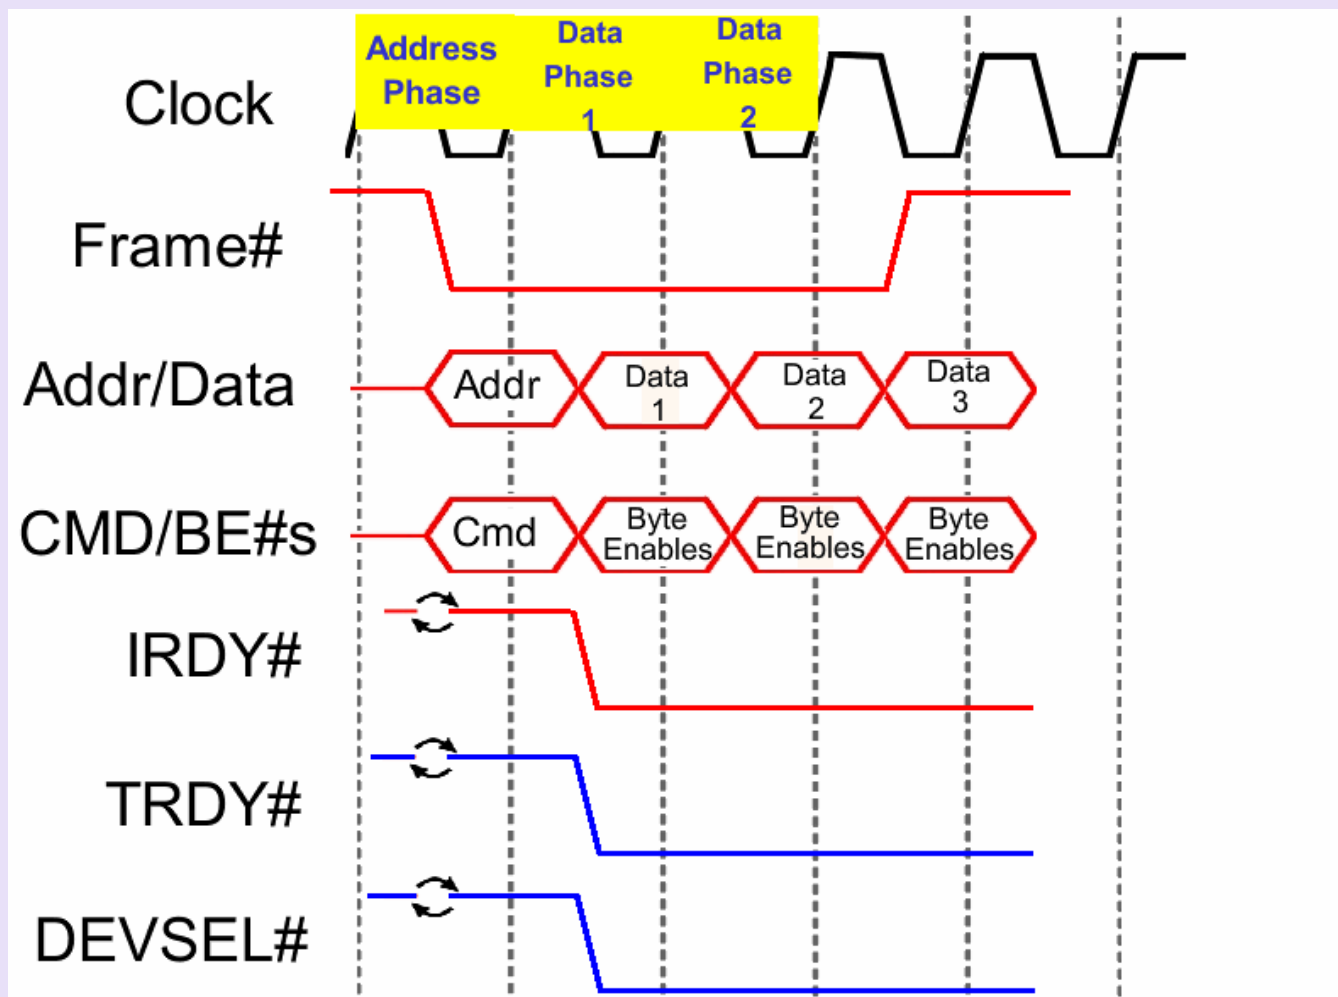


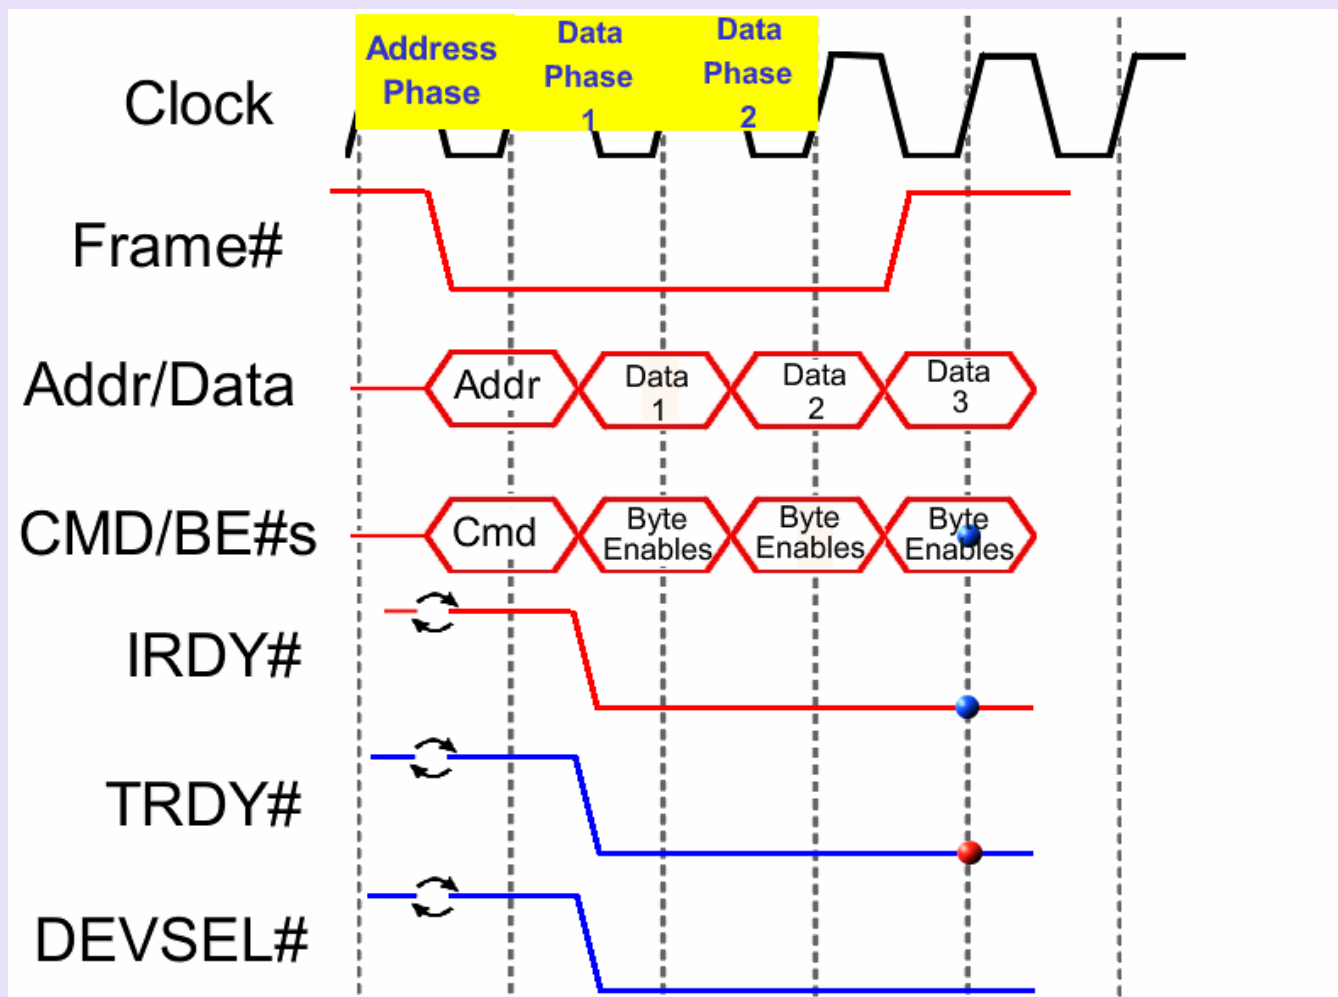
The Last Data Phase

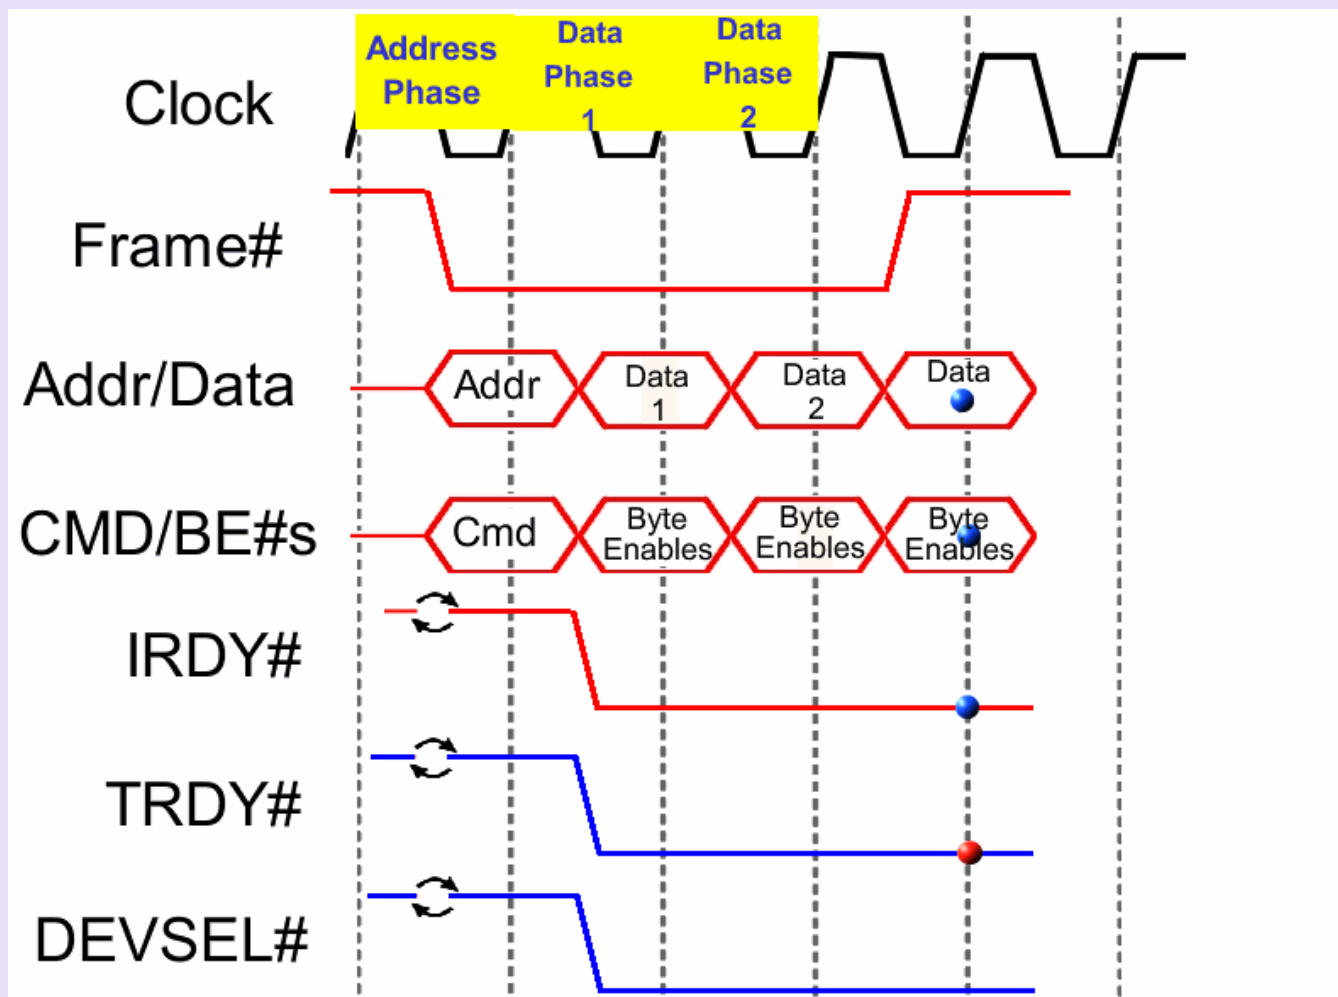


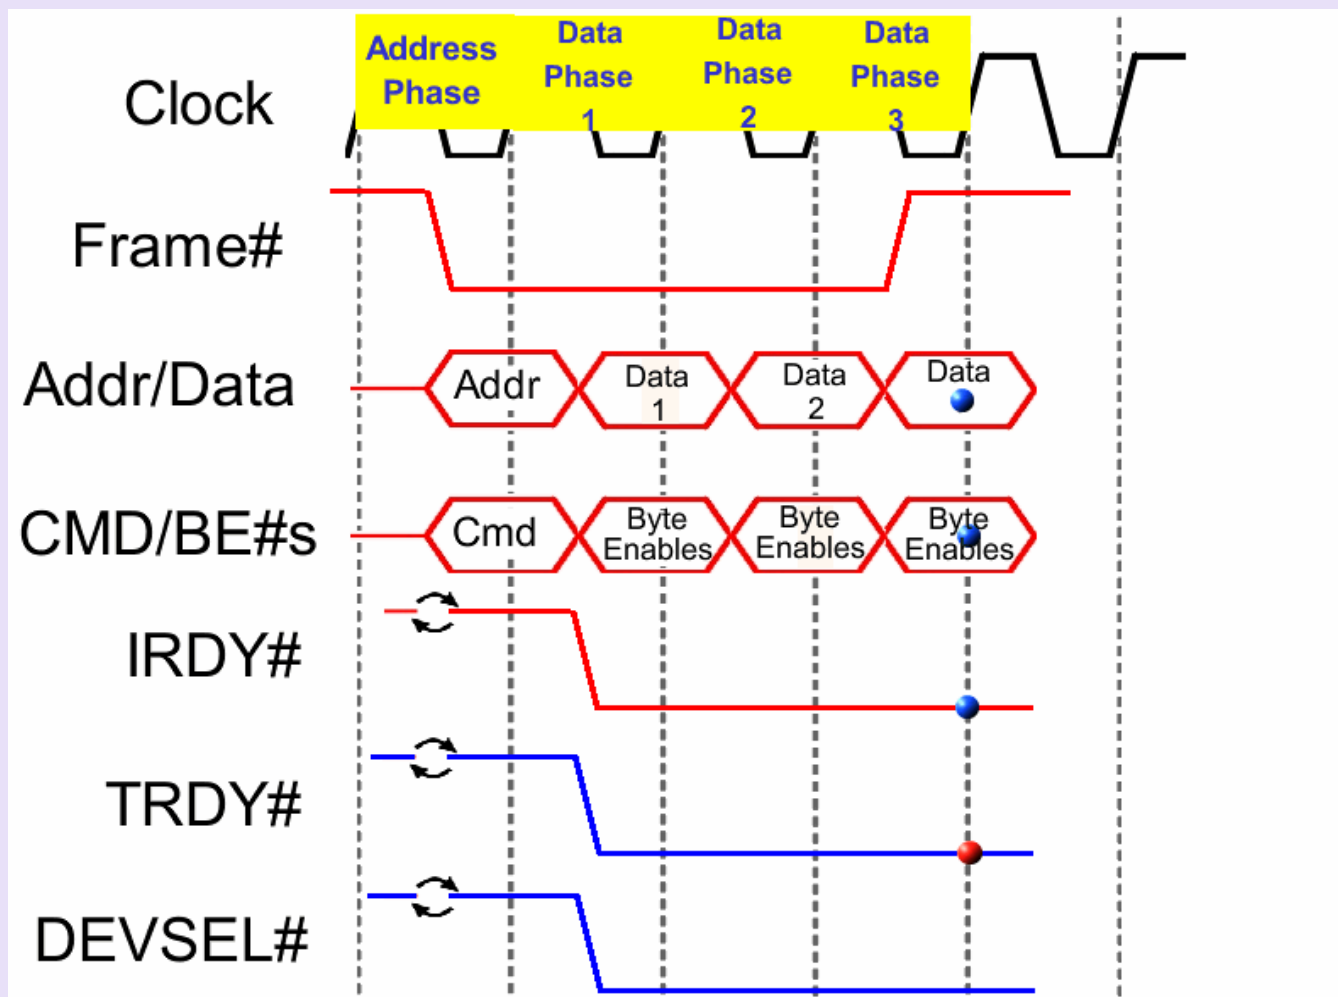





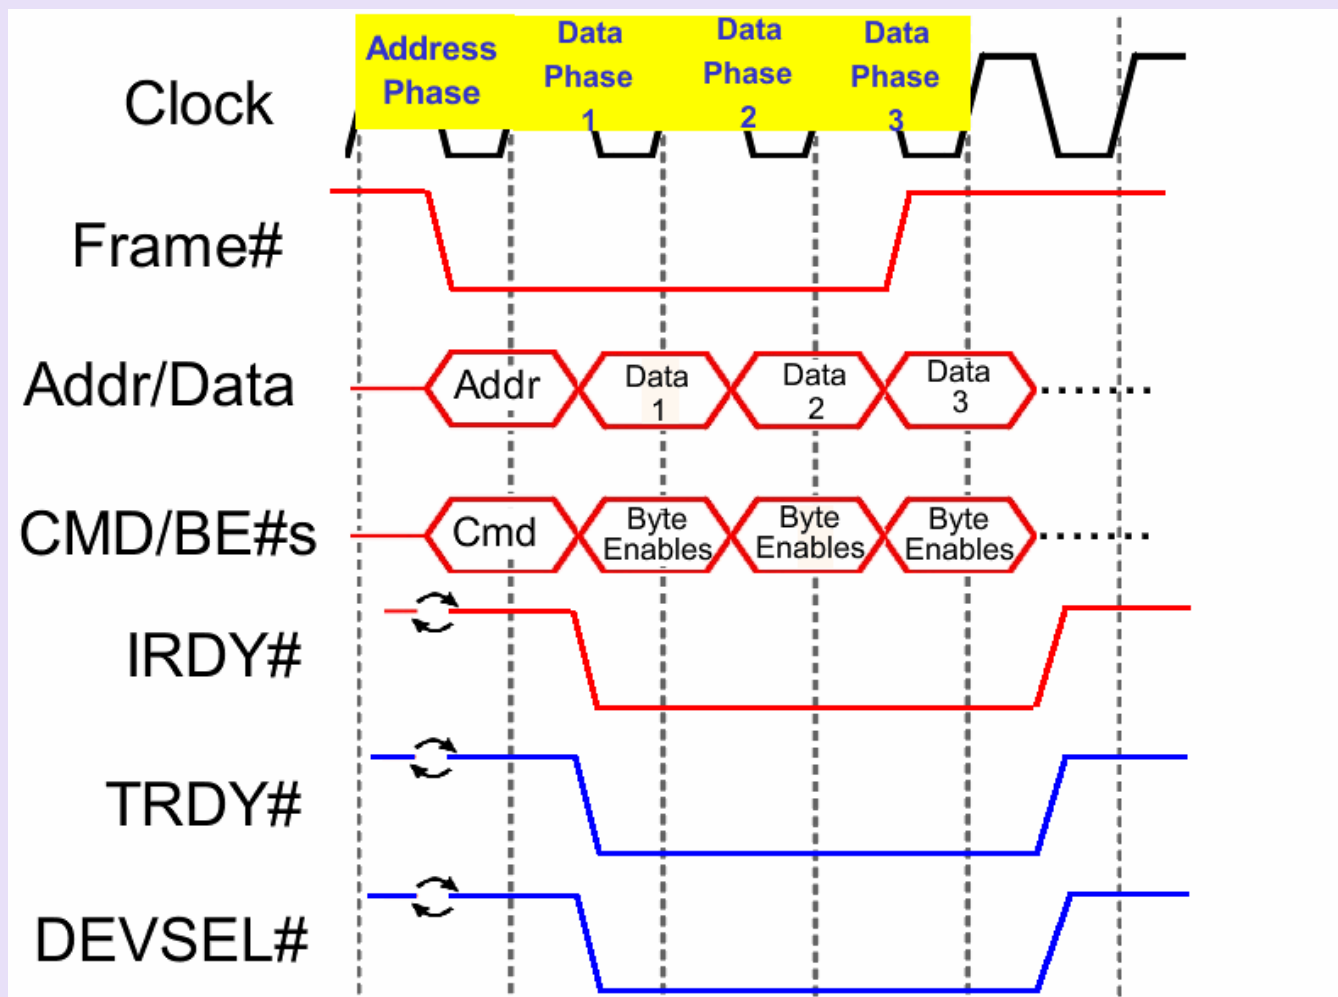












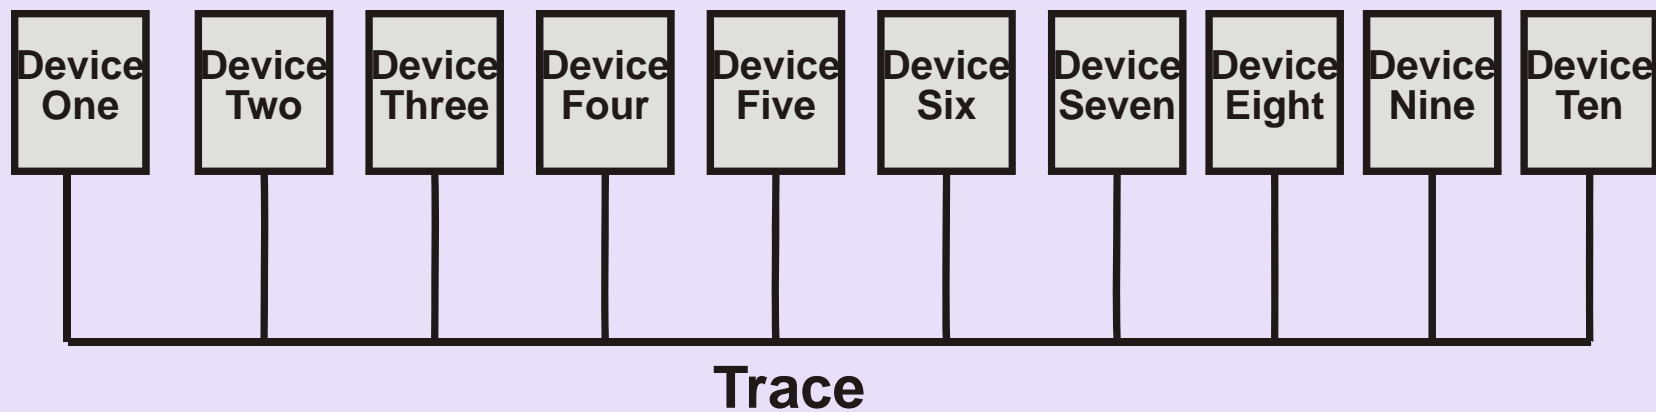





Memory Read Example




Reflected Wave Switching




Note that CLK is *not* a reflected-wave signal.




Note that CLK is *not* a reflected-wave signal.



Bus Master-related Signals

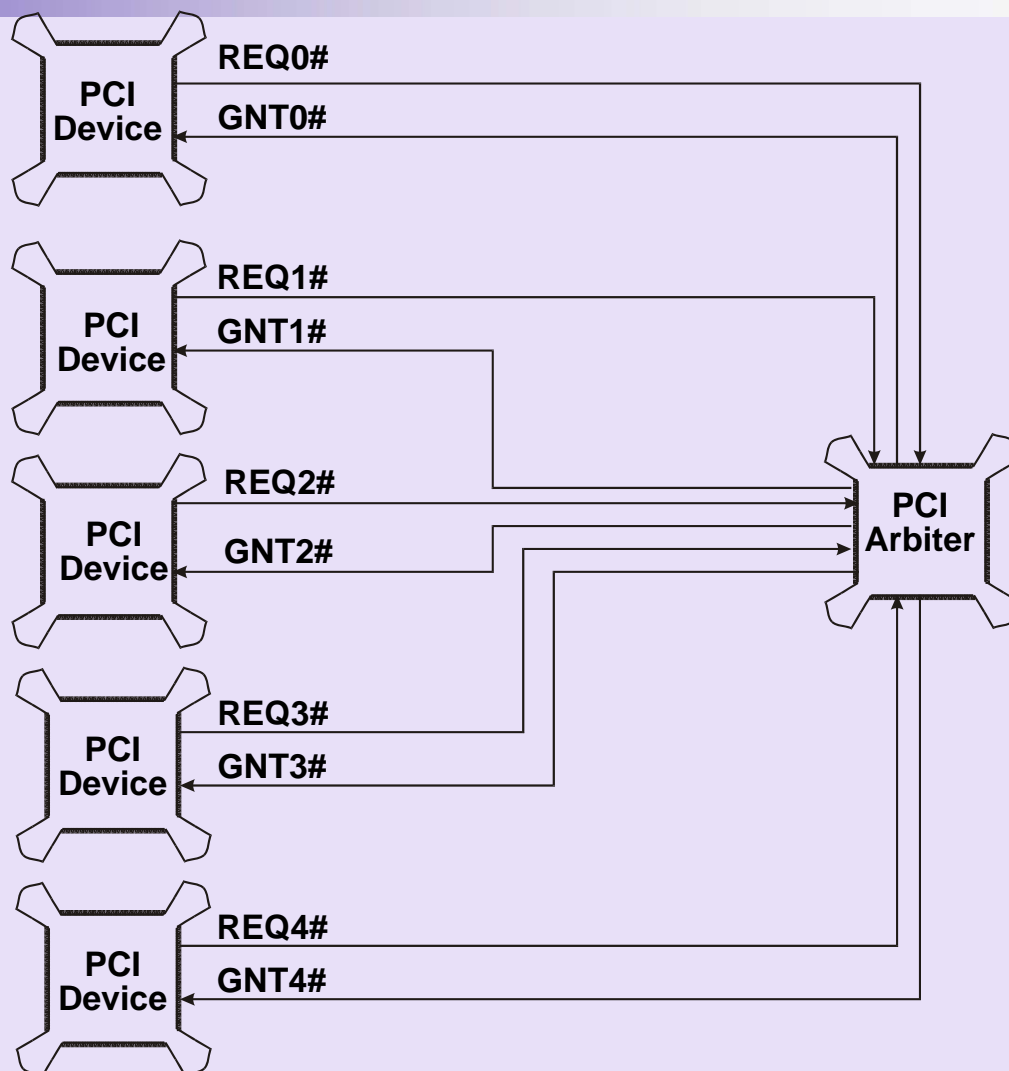


Target-related Signals



Arbitration Signaling


Hidden Bus Arbitration



Example Arbiter with fairness

First Group: Short Latency Devices

Second Group: Long Latency Devices

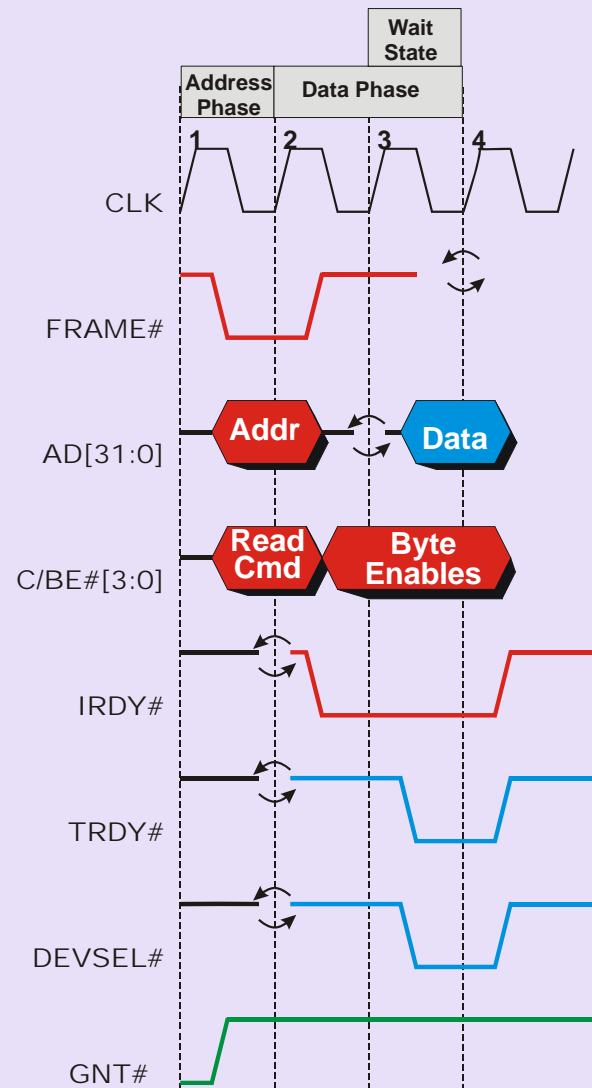


The Command Set


Table 1-1: PCI Command Types

C/BE[3:0]# (binary)	Command Type
0000	Interrupt Acknowledge
0001	Special Cycle
0010	I/O Read
0011	I/O Write
0100	Reserved
0101	Reserved
0110	Memory Read
0111	Memory Write
1000	Reserved
1001	Reserved
1010	Configuration Read
1011	Configuration Write
1100	Memory Read Multiple
1101	Dual Address Cycle
1110	Memory Read Line
1111	Memory Write-and-Invalidate

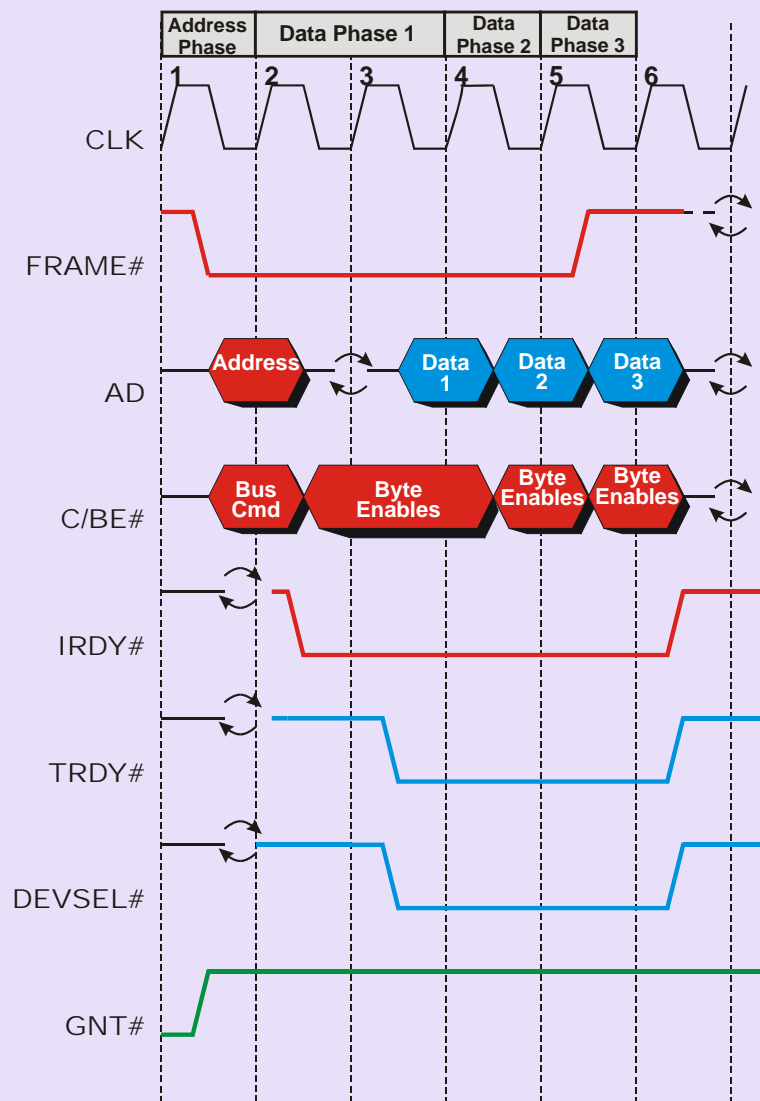
Single Data Phase Read



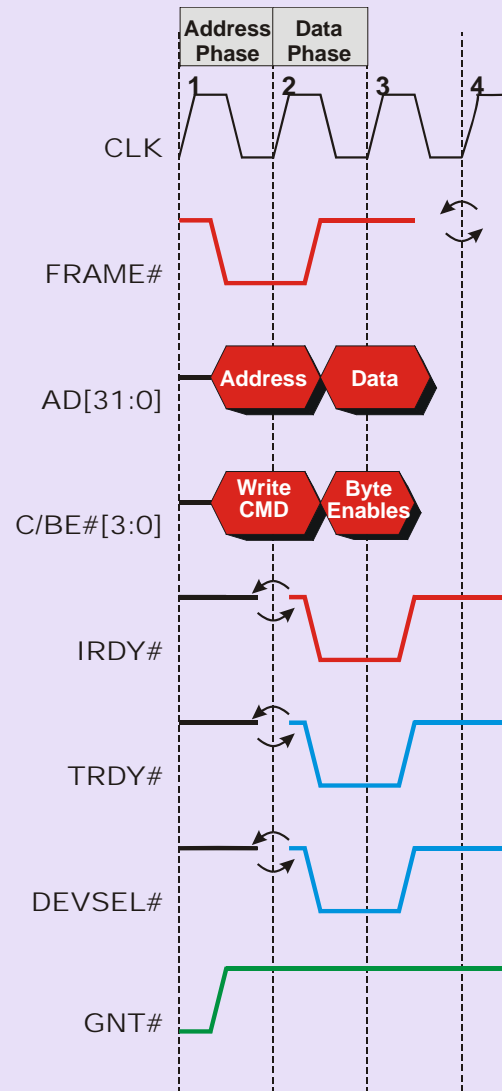
Example Burst Read



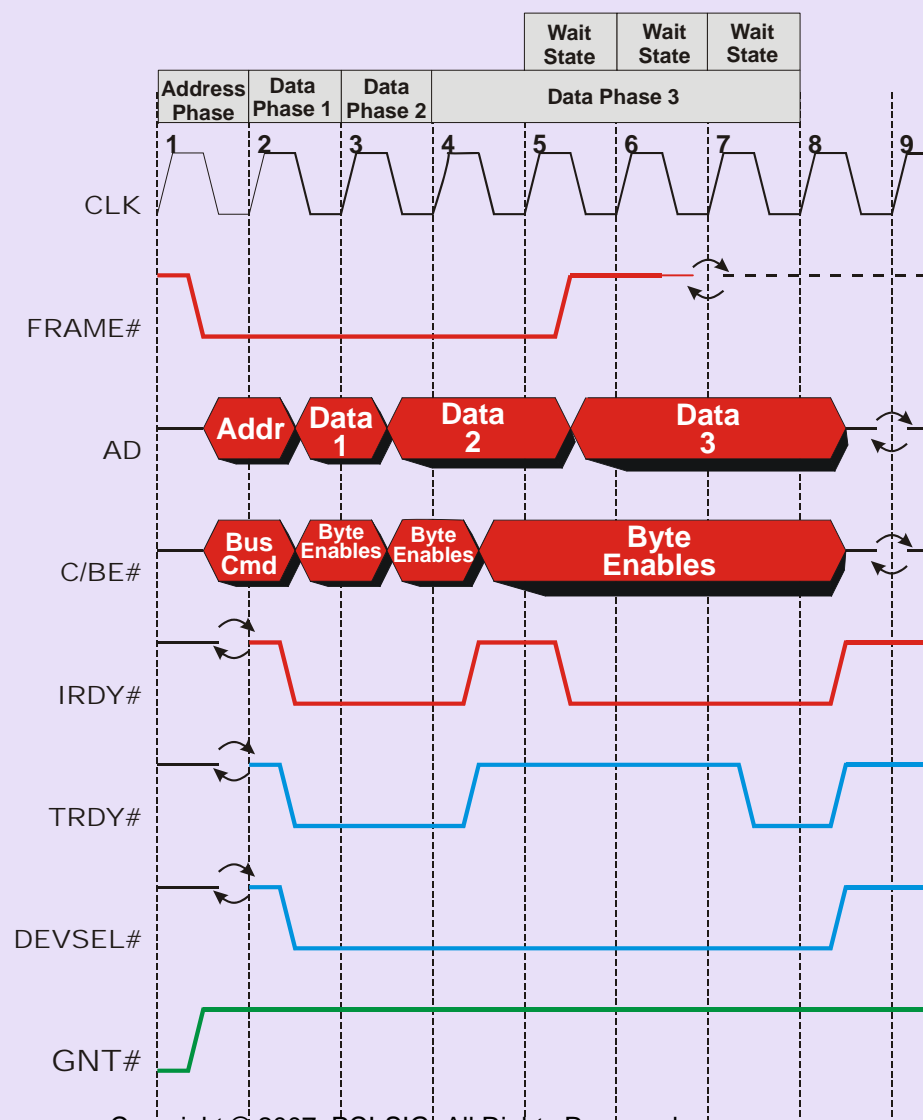
Optimized Burst Read



Single Data Phase Write

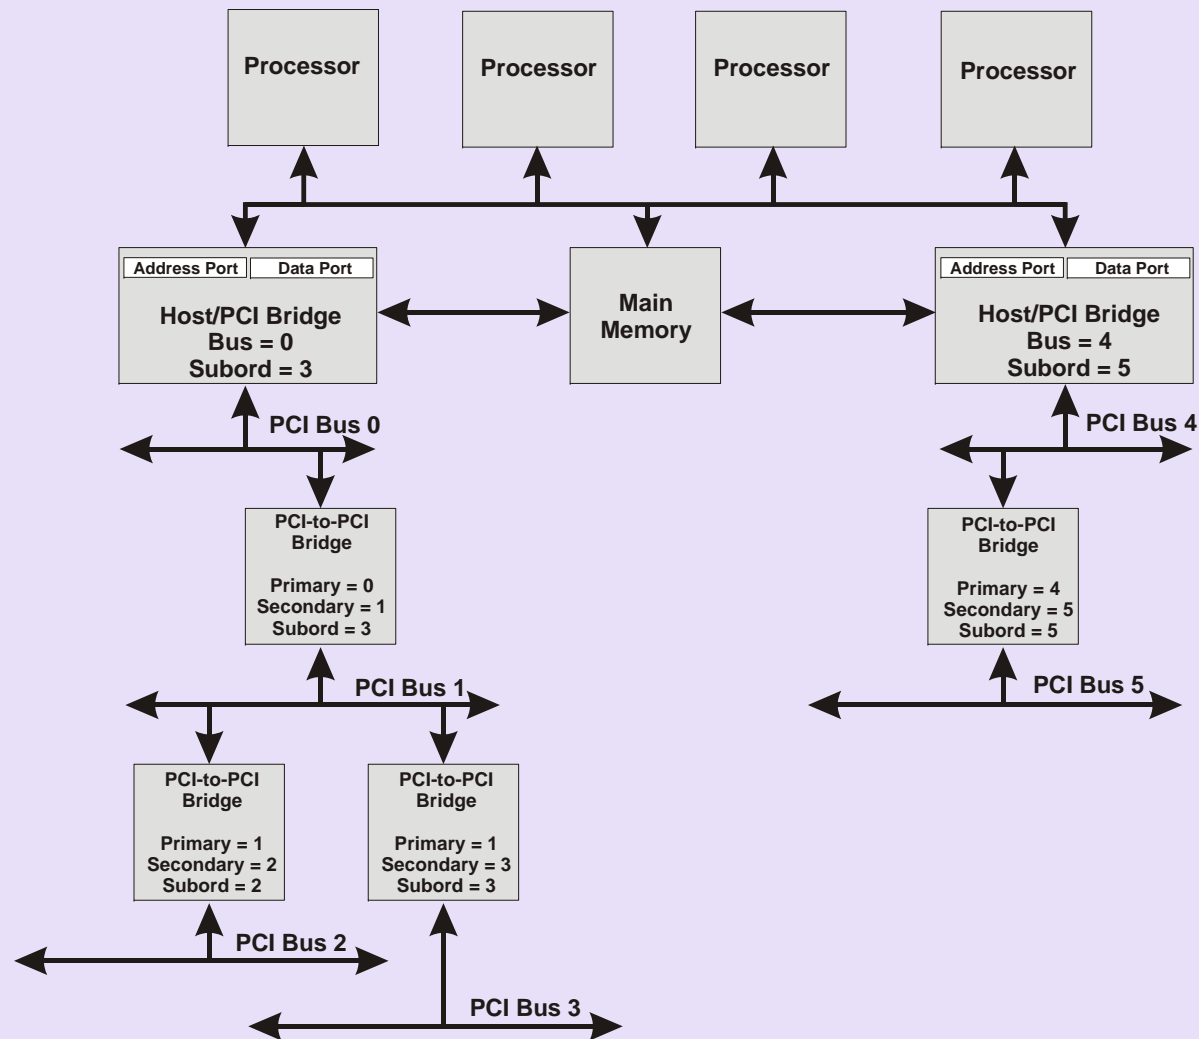


Burst Write



Configuration Transactions

- Each PCI device must be identified and configured.
- Devices are identified by:
 - ✓ Bus number on which they reside
 - ✓ Device number
 - ✓ Function number within the device
 - ✓ Byte location with 256 Bytes of config. space



Configuration Address Generation (mechanism #1)

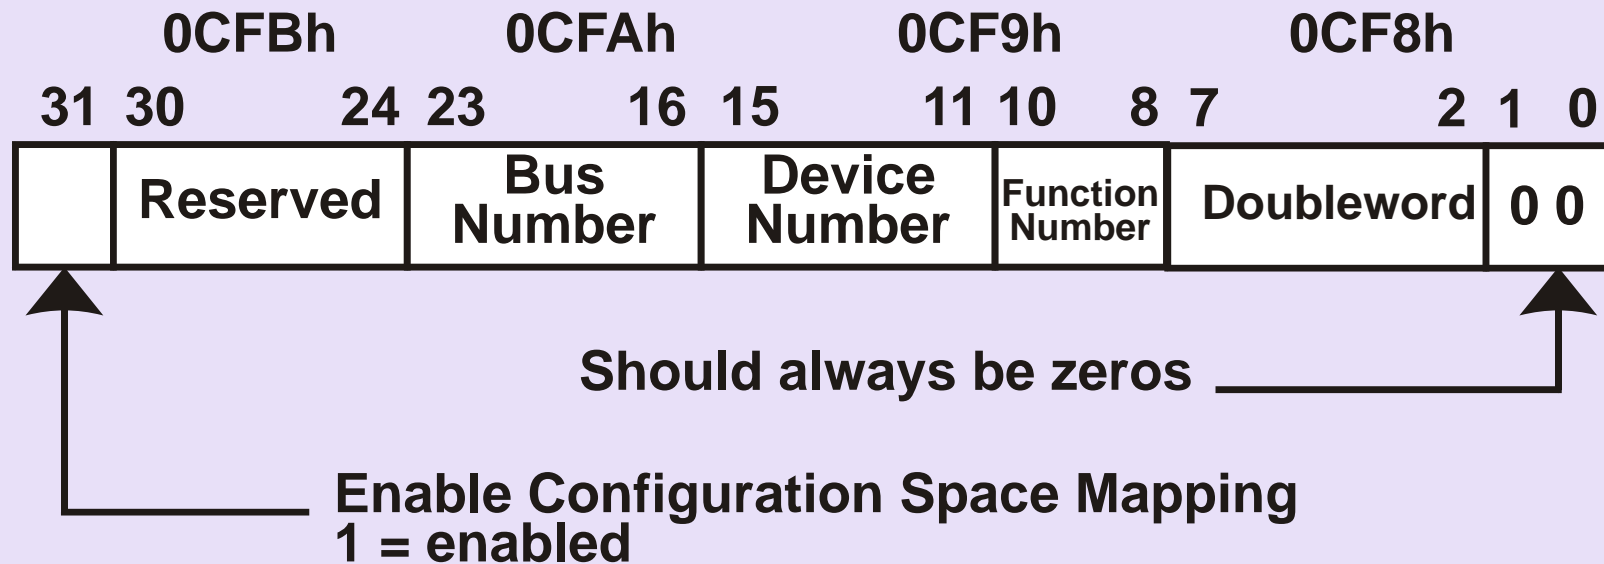
Configuration Address Register

Writes to the Configuration Address Register
Identifies configuration location to be accessed.

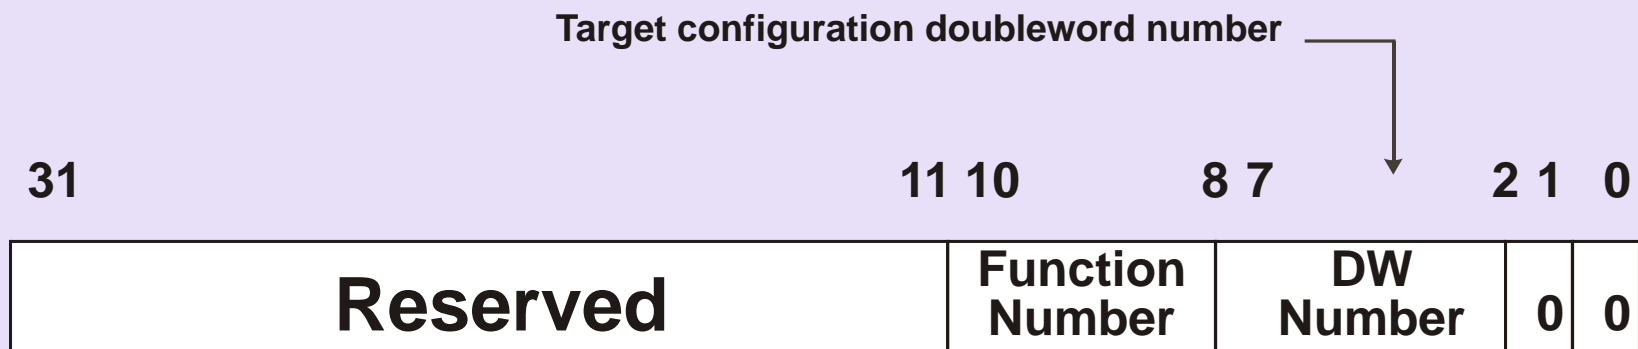
Configuration Data Register

Reads from or Writes to the Configuration Data Register
Causes a configuration read or write transaction to be Performed.


Configuration Address Registers




Config. Cycle Address Bus Contents



IDSEL Generation




IDSEL Generation

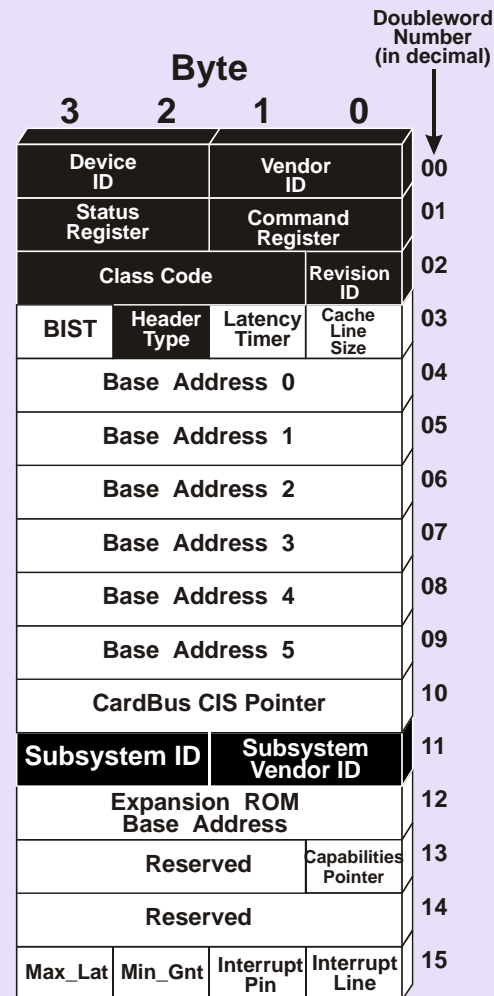


Type 1 Configuration Transactions

Doubleword number in device's configuration space




Header Type Zero




■ Required configuration registers

Registers that ID the Device

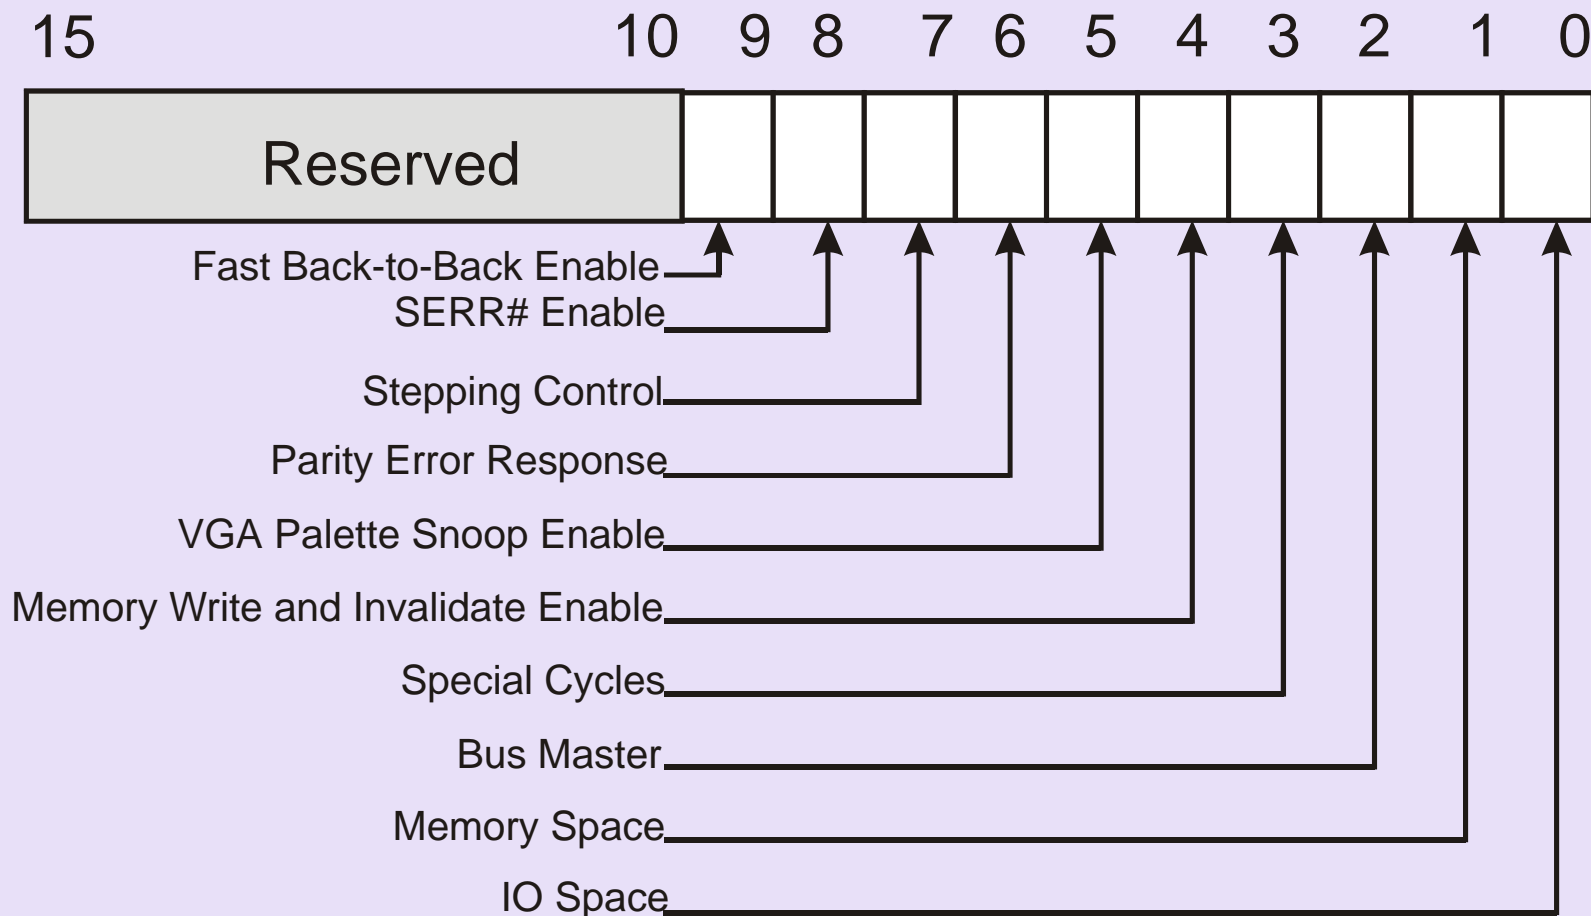


Required configuration registers

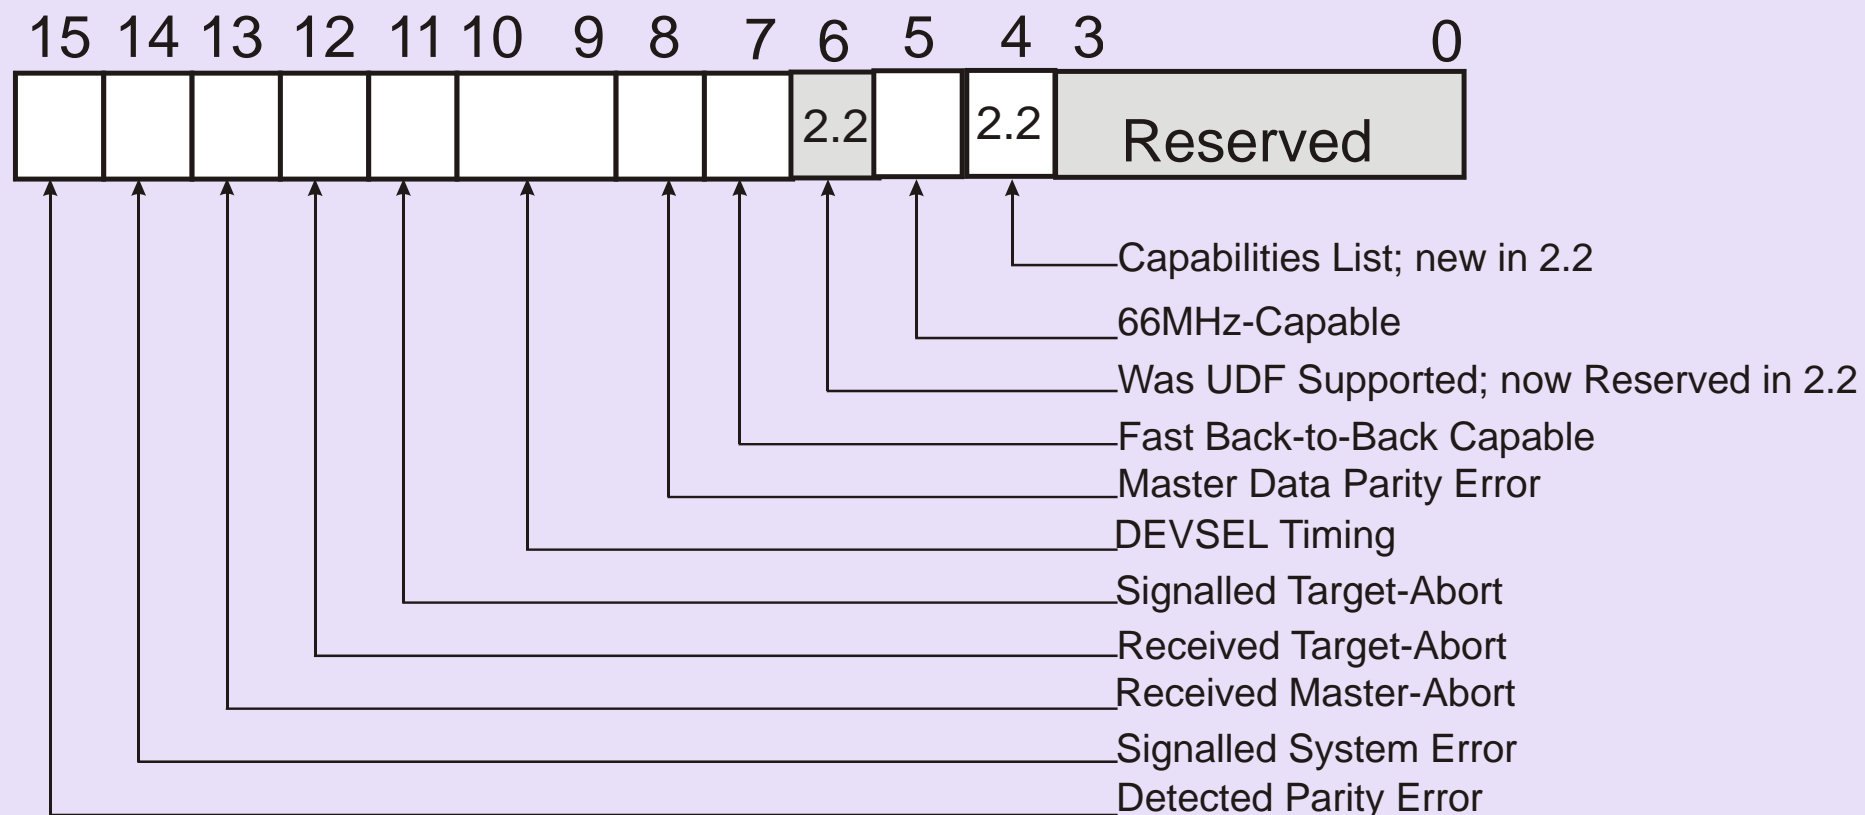
Class Code Register Format




Command Register Format




Status Register Format



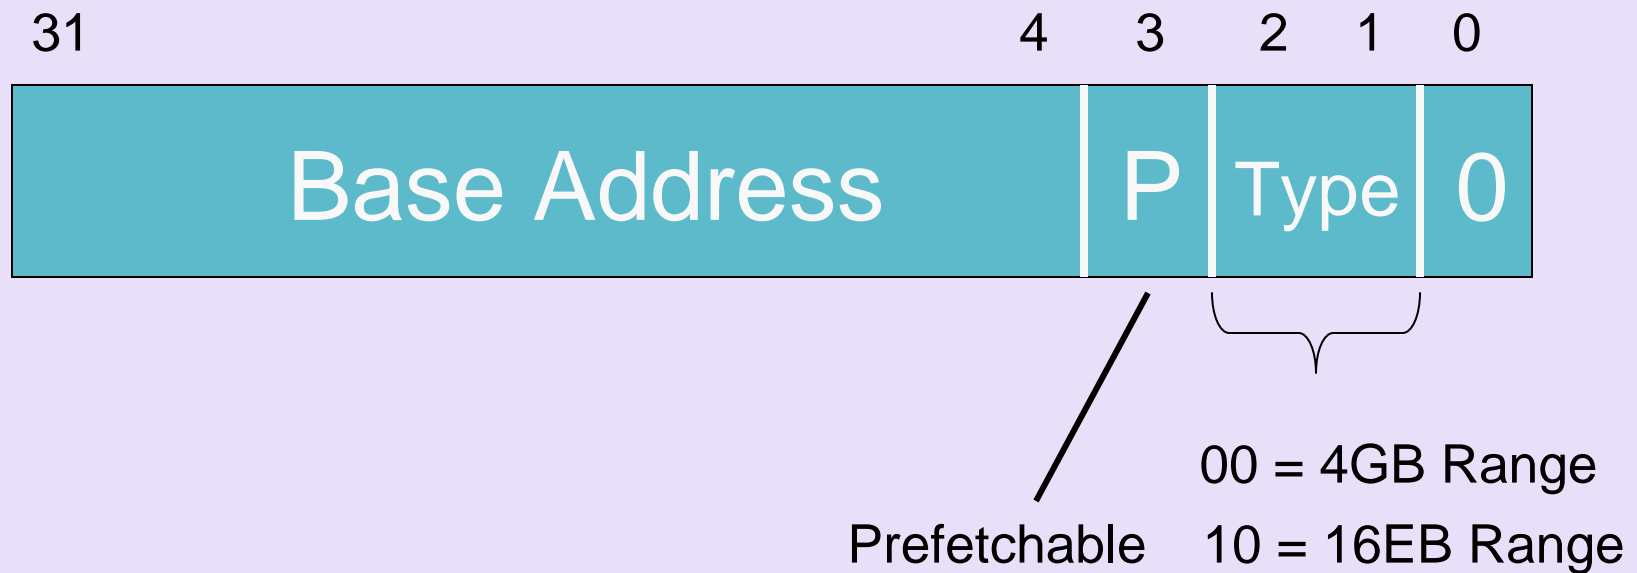
Header Type Register Format



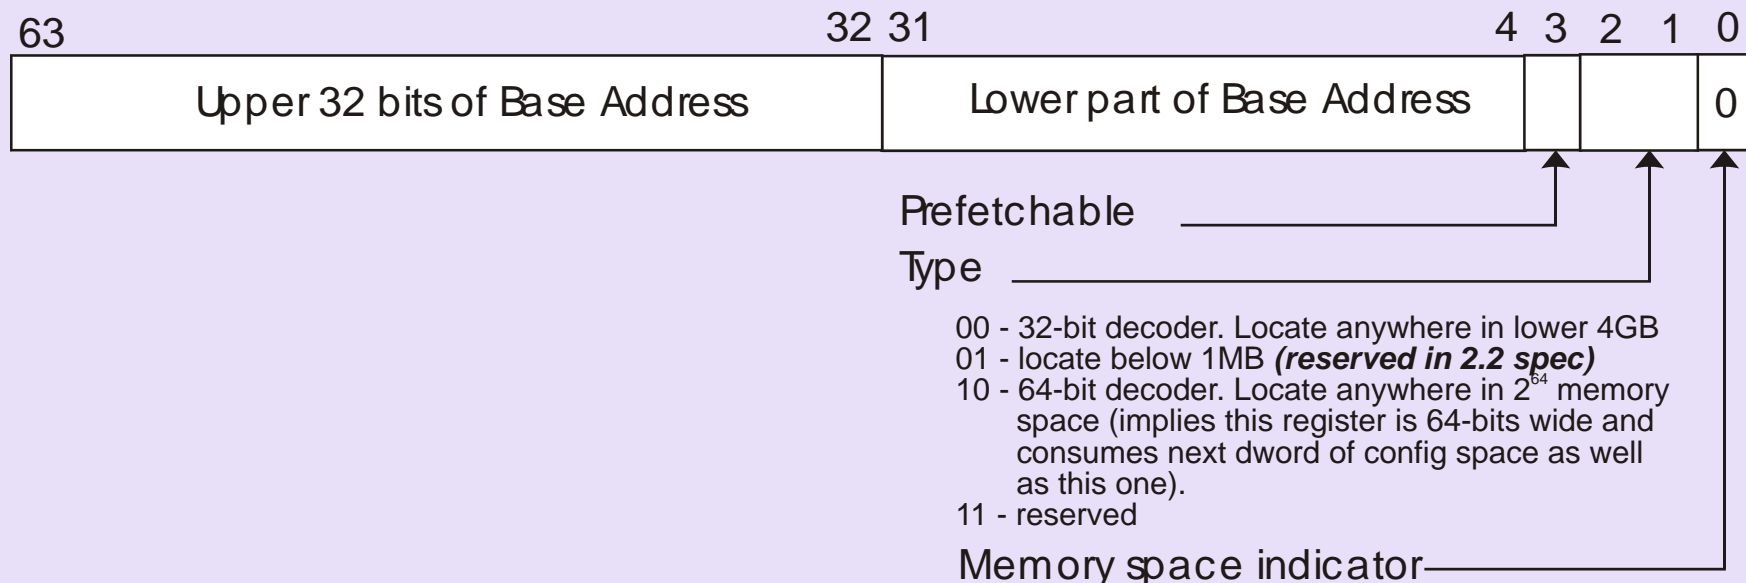
BIST Register Format



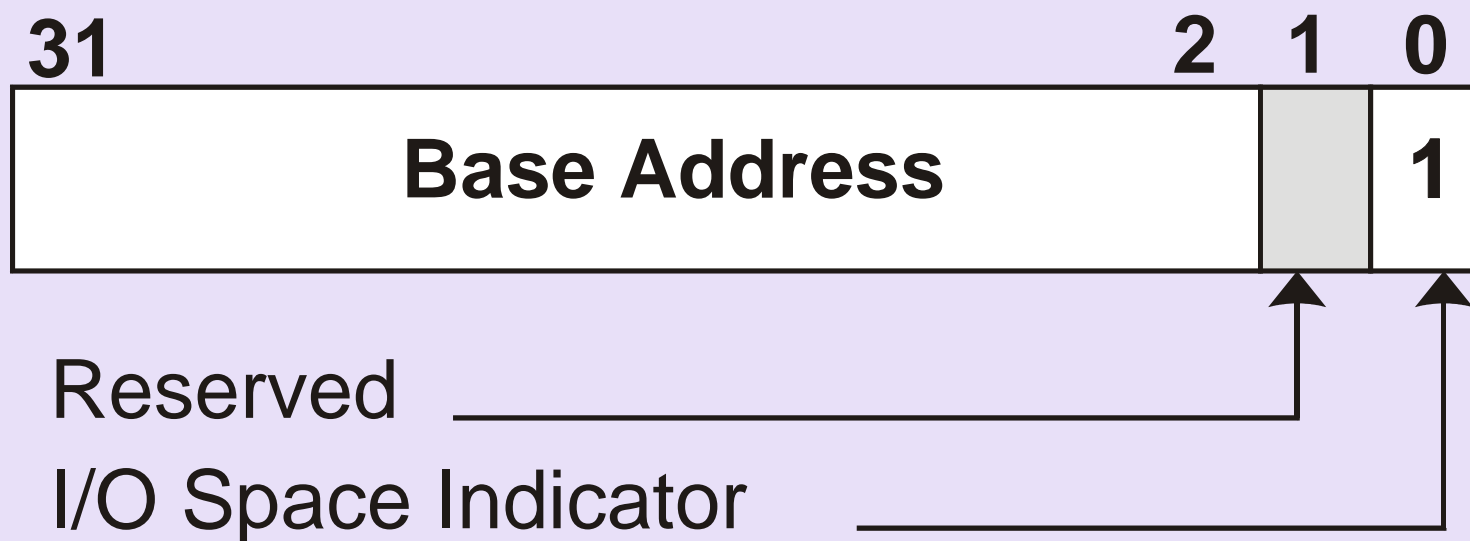
32-Bit Memory Bar




64-Bit Memory Bar



IO Bar



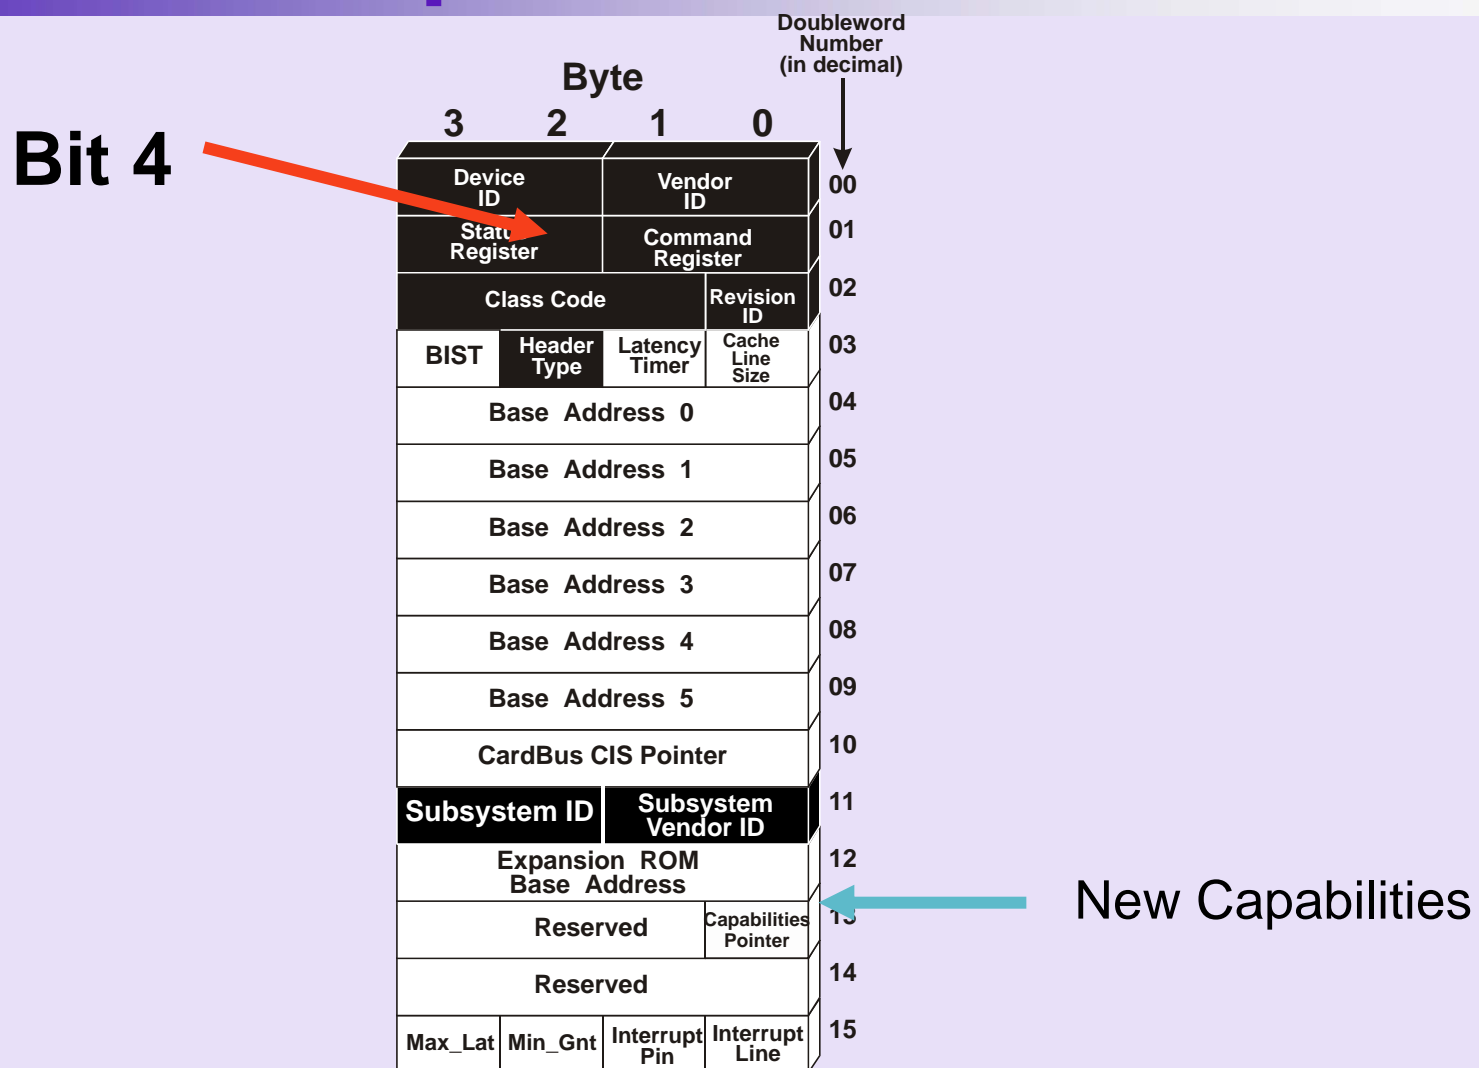
ROM Bar



Address decode enable _____



New Capabilities



Capability IDs

00h	Reserved
01h	PCI Power Management
02h	AGP
03h	VPD
04h	Slot Identification
05h	Message Signaled Interrupts - MSI
06h	CompactPCI Hot Swap
07h	PCI-X
10h	PCI Express
11h	Enhanced Message Signaled Interrupts - MSI-X

Capability Register Format



PCI Interrupt Deliver Mechanisms

- INTx Interrupt Delivery (x=A, B, C, or D)
- Message Signaled Interrupts - MSI
(via Memory Writes)
- Enhanced Message Signaled Interrupts - MSI-X
(via Memory Writes)


Function Reports Interrupt Pin

00h = no interrupts
 01h = INTA#
 02h = INTB#
 03h = INTC#
 04h = INTD#

Byte				Dword Number (in decimal)
3	2	1	0	
Device ID		Vendor ID		00
Status Register		Command Register		01
Class Code			Revision ID	02
BIST	Header Type	Latency Timer	Cache Line Size	03
Base Address 0				04
Base Address 1				05
Base Address 2				06
Base Address 3				07
Base Address 4				08
Base Address 5				09
CardBus CIS Pointer				10
Subsystem ID		Subsystem Vendor ID		11
Expansion ROM Base Address				12
Reserved			Capabilities Pointer	13
Reserved				14
Max_Lat	Min_Gnt	Interrupt Pin	Interrupt Line	15

■ Required configuration registers

Interrupt Signal Routing



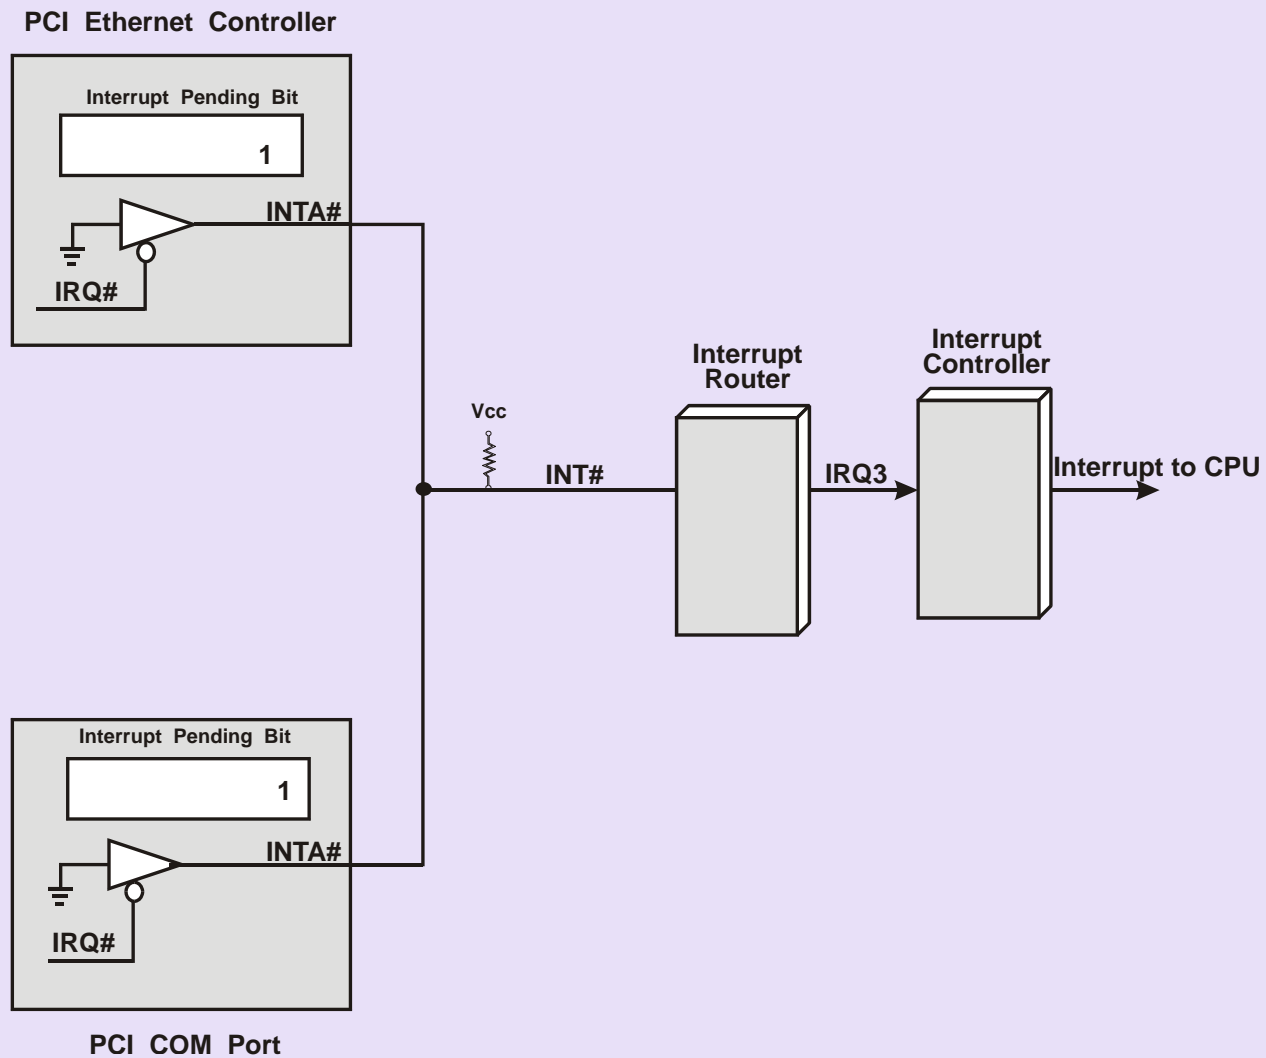
Interrupt Line Register

0 = IRQ0
 1 = IRQ1
 2 = IRQ2
 3 = IRQ3
 4 = IRQ4
 5 = IRQ5
 6 = IRQ6
 7 = IRQ7
 8 = IRQ8
 9 = IRQ9
 A = IRQ10
 B = IRQ11
 C = IRQ12
 D = IRQ13
 E = IRQ14
 F = IRQ15

Byte				Doubleword Number (in decimal)
3	2	1	0	
Device ID		Vendor ID		00
Status Register		Command Register		01
Class Code			Revision ID	02
BIST	Header Type	Latency Timer	Cache Line Size	03
Base Address 0				04
Base Address 1				05
Base Address 2				06
Base Address 3				07
Base Address 4				08
Base Address 5				09
CardBus CIS Pointer				10
Subsystem ID		Subsystem Vendor ID		11
Expansion ROM Base Address				12
Reserved			Capabilities Pointer	13
Reserved				14
Max_Lat	Min_Gnt	Interrupt Pin	Interrupt Line	15

Required configuration registers

Interrupt Sharing



Message Signaled Interrupts

- Allow master to send interrupts via Memory Write transaction.
- Software assigns memory address location to write to within CPU/PCI bridge and assigns data value to be written.

Advantages of MSI

- No interrupt traces required
- No interrupt sharing
- No device driver chaining
- MSI can identify the service to be performed
- No read needed to ensure write buffers are flushed



MSI Capability Register 32 & 64 Bit Versions



31	16	15	8	7	0	Offset
Message Control Register			Pointer to Next ID		Capability ID = 05h	00h
Message Address Register						04h
Reserved			Message Data Register			08h
Mask Bits						0Ch
Pending Bits						10h

31	16	15	8	7	0	
Message Control Register			Pointer to Next ID		Capability ID = 05h	00h
Least-Significant 32 bits of Message Address Register						04h
Most-Significant 32 bits of Message Address Register						08h
Reserved			Message Data Register			0Ch
Mask Bits						10h
Pending Bits						14h

What do you get with MSI-X?


- An MSI-X capable device can deliver interrupts to any Processor in an SMP platform
- Limit of 32 vectors per device eliminated
- Each vector can target a different processor
 - ✓ Accomplished by supporting a separate Address/Data register per vector and not relying on re-vectoring table implemented in chipset

MSI-X Capability Register

31	16	15	8	7	2	0	Offset											
Message Control Register																Pointer to Next ID	Capability ID = 11h	00h
MSI-X Table Offset																Table BIR		04h
Pending Bit Array (PBA) Offset																PBA BIR		08h

BIR = BAR Indicator Register

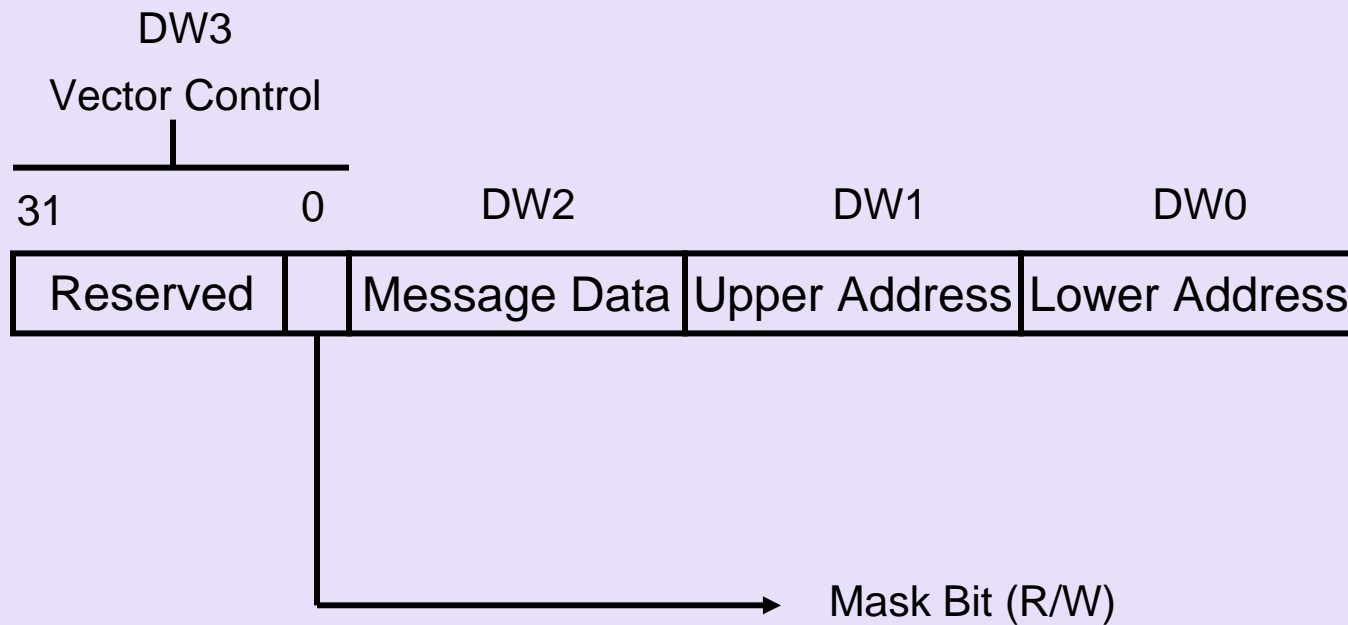
MSI-X Control Register



MSI-X Table Structure

DW3	DW2	DW1	DW0	
Vector Control	Message Data	Upper Address	Lower Address	Entry 0
Vector Control	Message Data	Upper Address	Lower Address	Entry 1
Vector Control	Message Data	Upper Address	Lower Address	Entry 2
....	
....	
Vector Control	Message Data	Upper Address	Lower Address	Entry N-1

MSI-X Vector Control



MSI-X PBA Structure

DW1	DW0
Pending Bits 0 - 63	QW 0
Pending Bits 64 - 127	QW 1
Pending Bits 128 - 191	Entry 2
....	
....	
Pending Bits	QW (N-1)/64

Thank you for attending the PCI-SIG Developers Conference 2007

For more information please go to
www.pcisig.com



Conventional PCI Overview

Ravi Budruk
Senior Staff Engineer and Partner
MindShare, Inc.