Introducing...

# QuickLogic's The Basics of PCI

*QuickPCI - The PCI Solution for System Needs*

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# Typical System Implementation

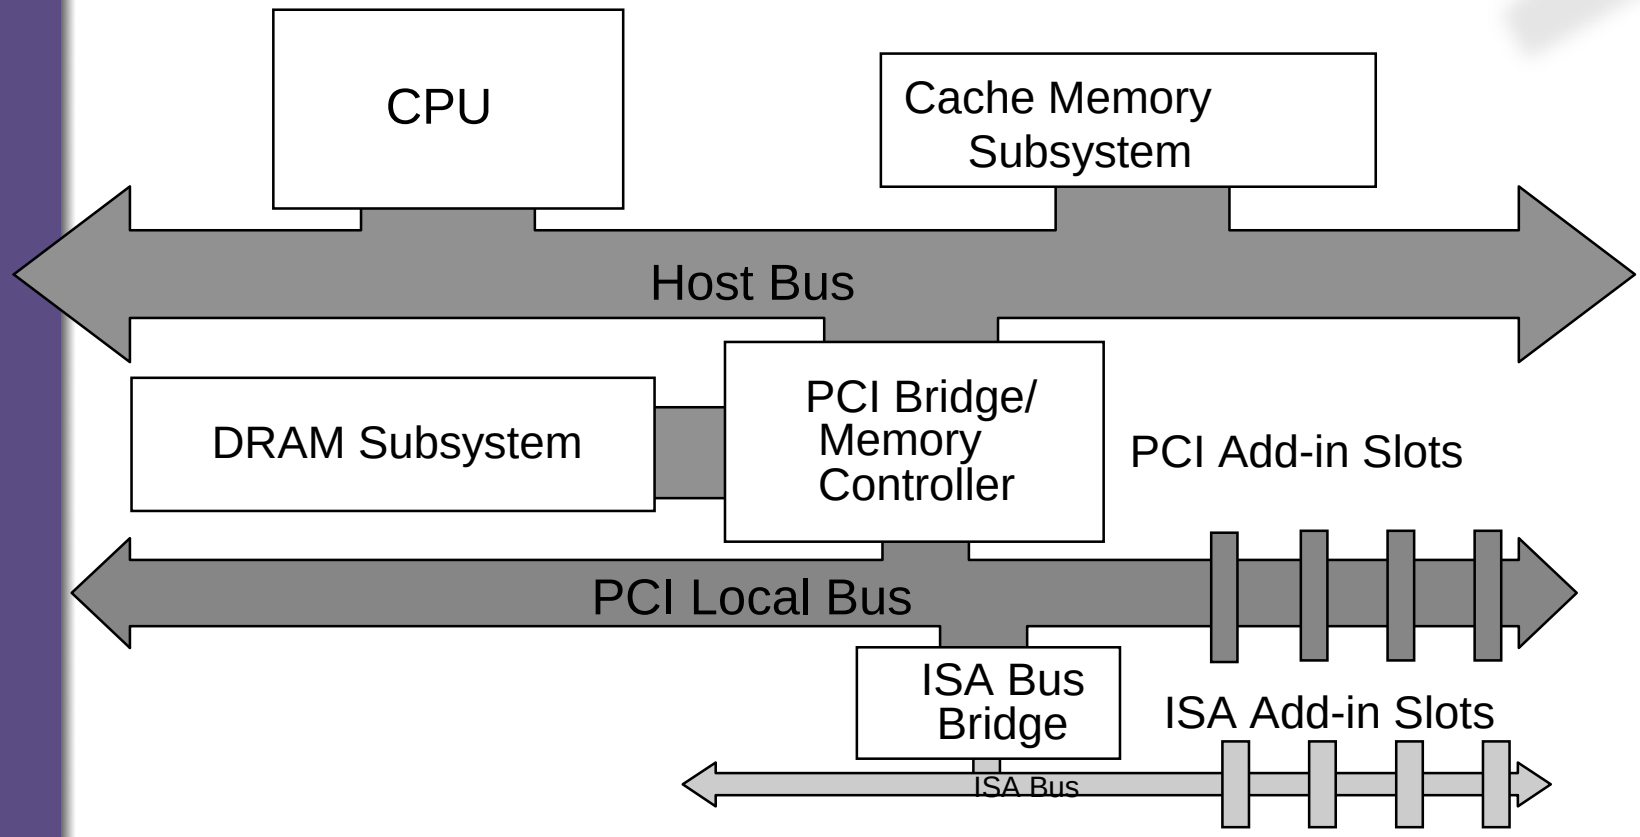

# Desktop Platforms

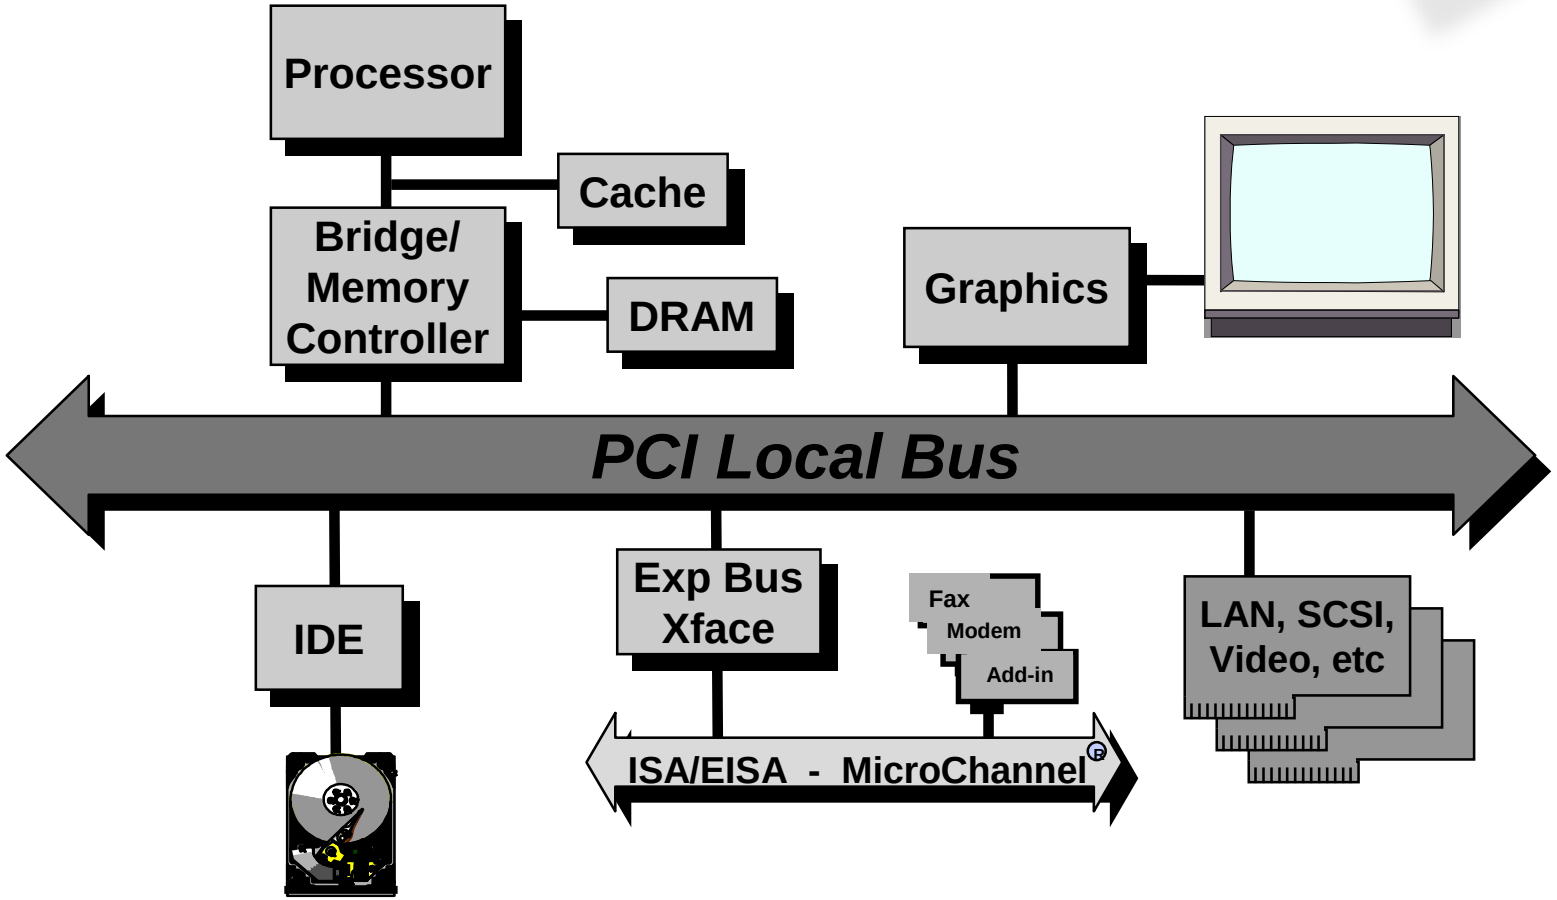

# Server Platforms

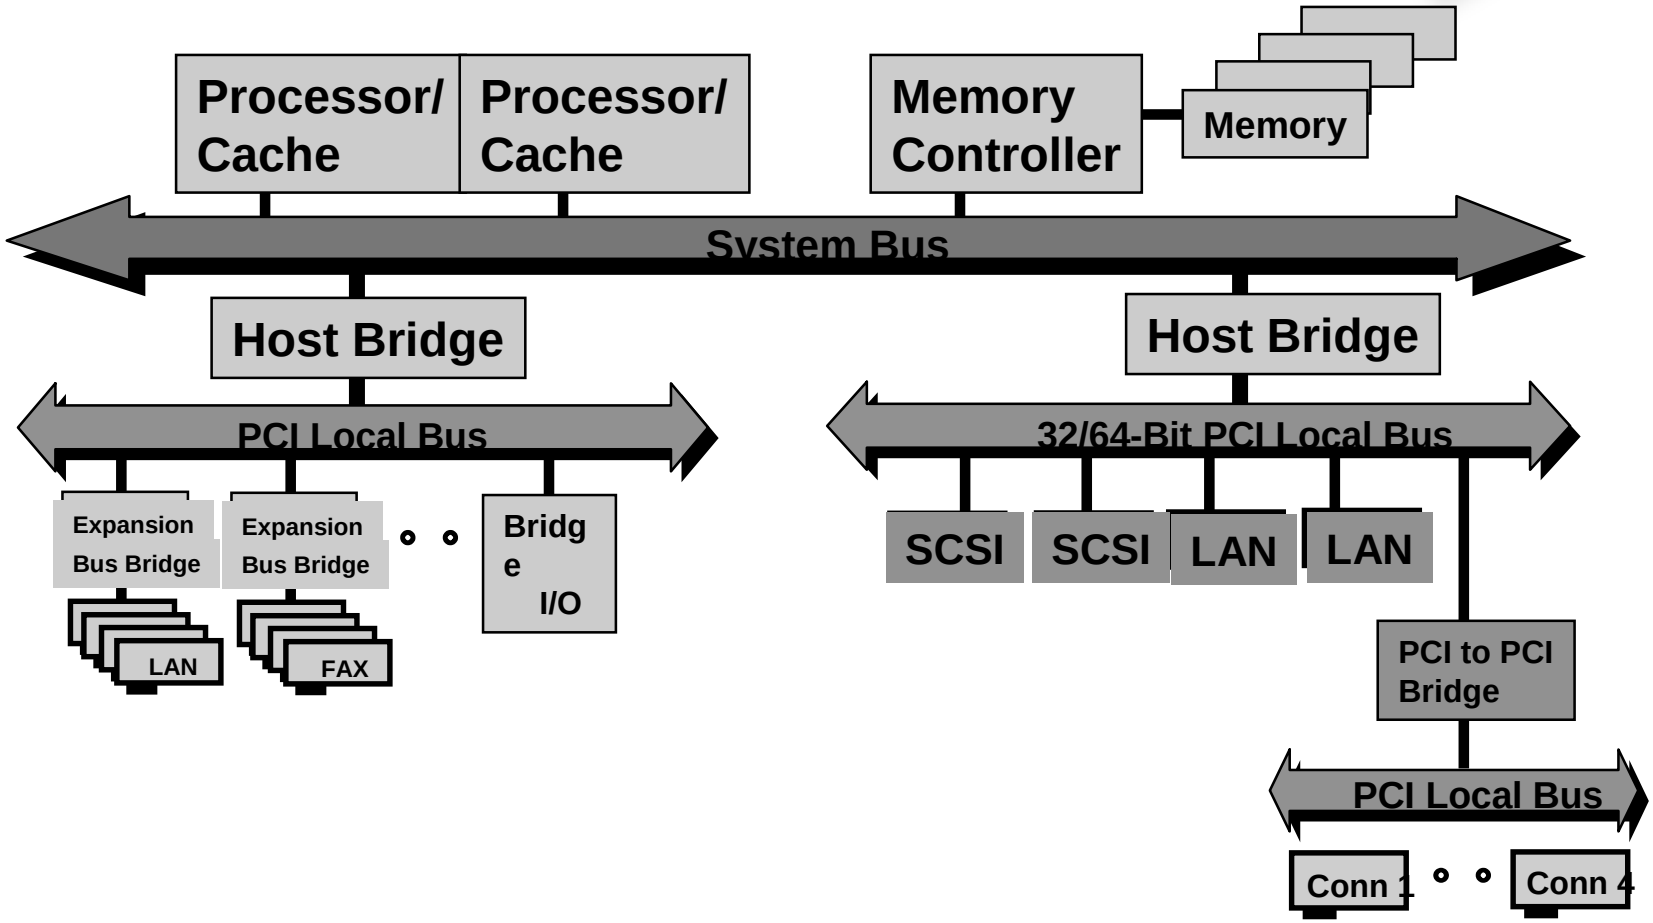

# Mobile Platforms

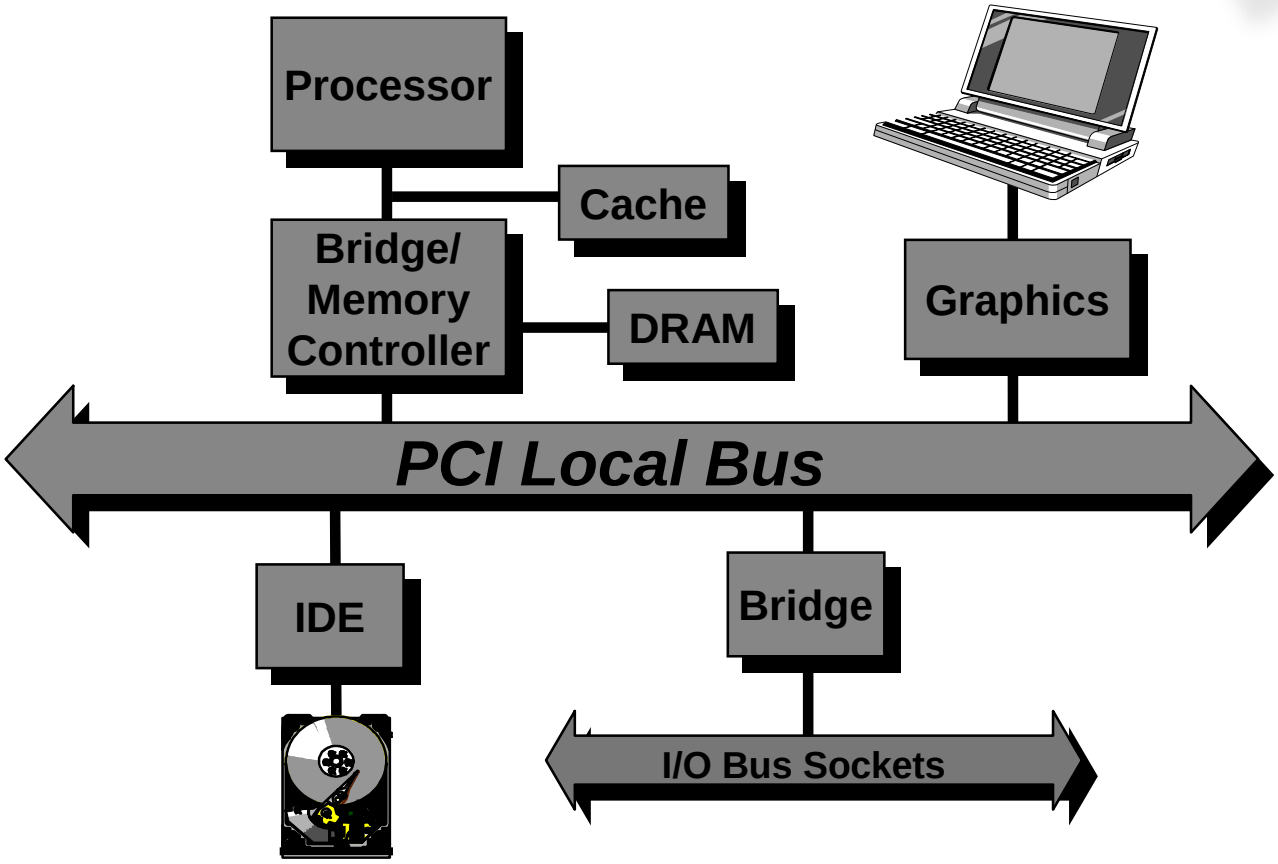

# PCI Auto-Configuration

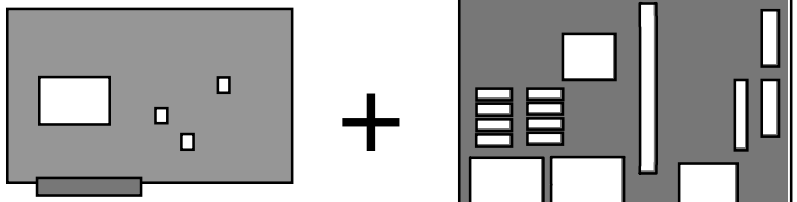

**PCI Add-in Card**

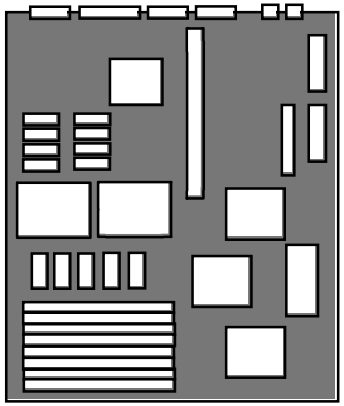

**PCI Based Mother Board**

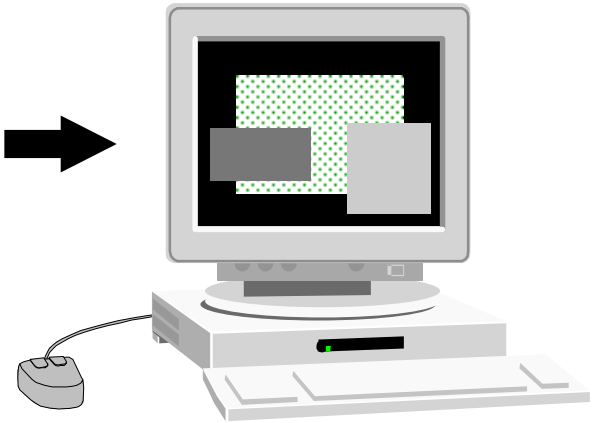

***Auto Configuration System***

# PCI Overview

# PCI Features

- 32-Bit or 64-Bit address and data
- 66 or 33 down to 0 MHz synchronous operation
- Single or multiple bus masters
- Reflected bus signaling
- Stepped signaling
- Bus parity error reporting
- 5 or 3.3 volt operation
- Cache support
- JTAG testing

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Bus

- **Bus Signals**
- **Bus Commands**
- **Bus Transactions**
- **Arbitration**

# PCI Bus Signals

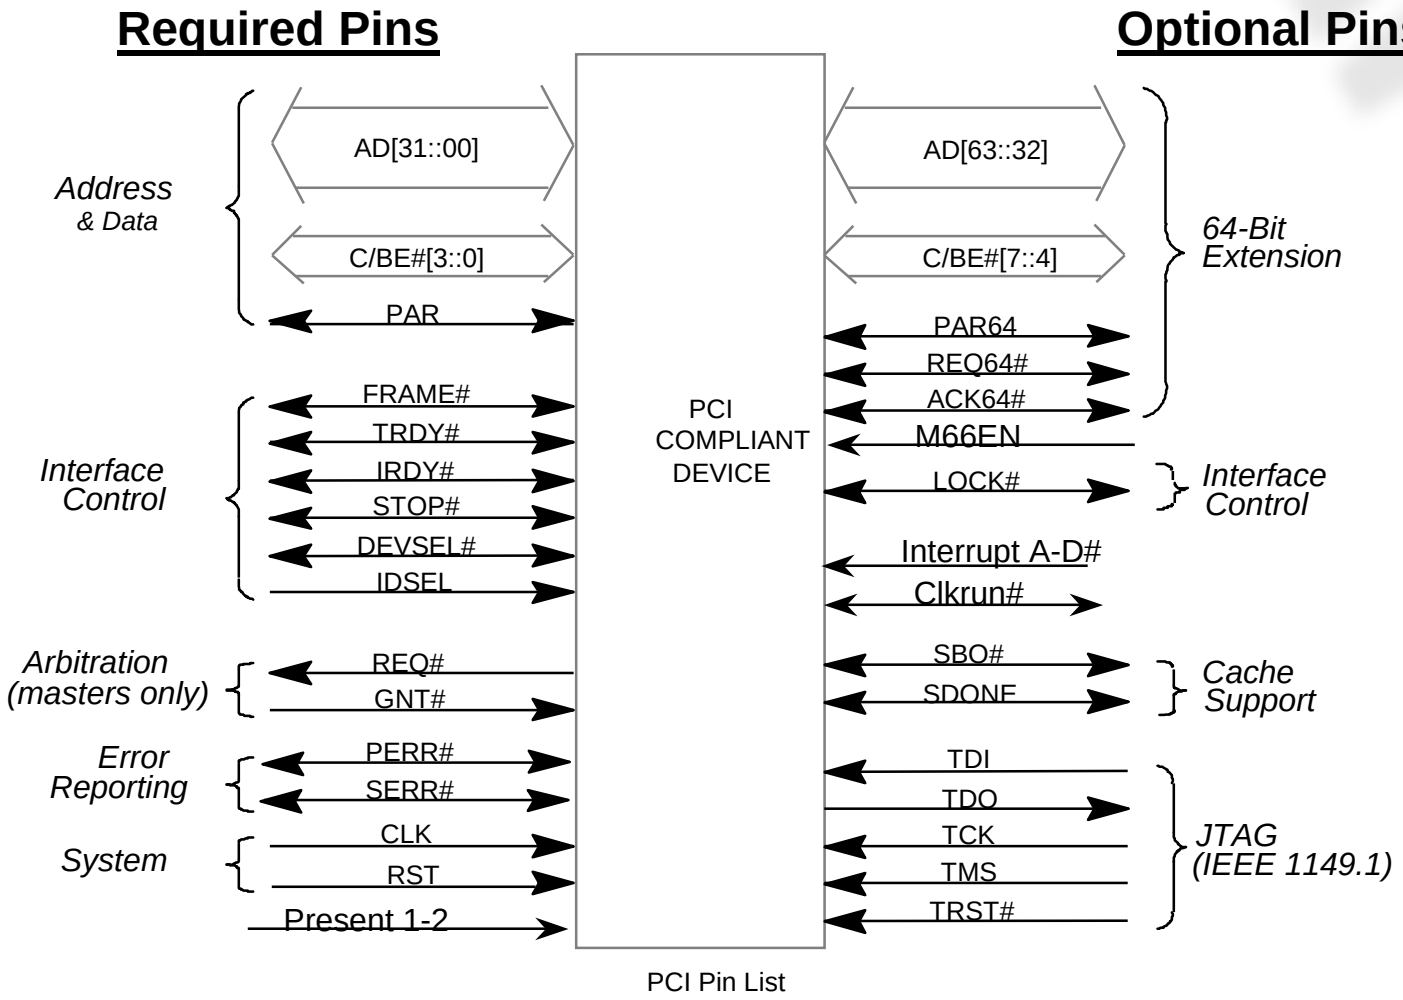


PCI Pin List

# PCI Bus Control Signals

- **System**
  - **CLK**
  - **RST#**
- **Address Data**
  - **AD[31:00]**
  - **C/BE[3:0]#**
  - **PAR**

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Bus Control Signals

- **Interface Control**
  - **FRAME#**
  - **IRDY#**
  - **TRDY#**
  - **STOP#**
  - **LOCK#**
  - **IDSEL#**
  - **DEVSEL#**
- **Arbitration**
  - **REQ#**
  - **GNT#**

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI bus access

- **PCI is a Multimaster Bus**
- **All transactions *initiated* by a master**
- **All transactions to/from a *target***

# PCI Bus Control Signals

- **FRAME#**
  - **driven by master to indicate transfer start and end**
- **IRDY#**
  - **driven by master to indicate it is ready to transfer data**
- **TRDY#**
  - **driven by master to indicate it is ready to transfer data**

QUICKLOGIC
Embedded Standard Products . . . a ger

# Bus transaction start

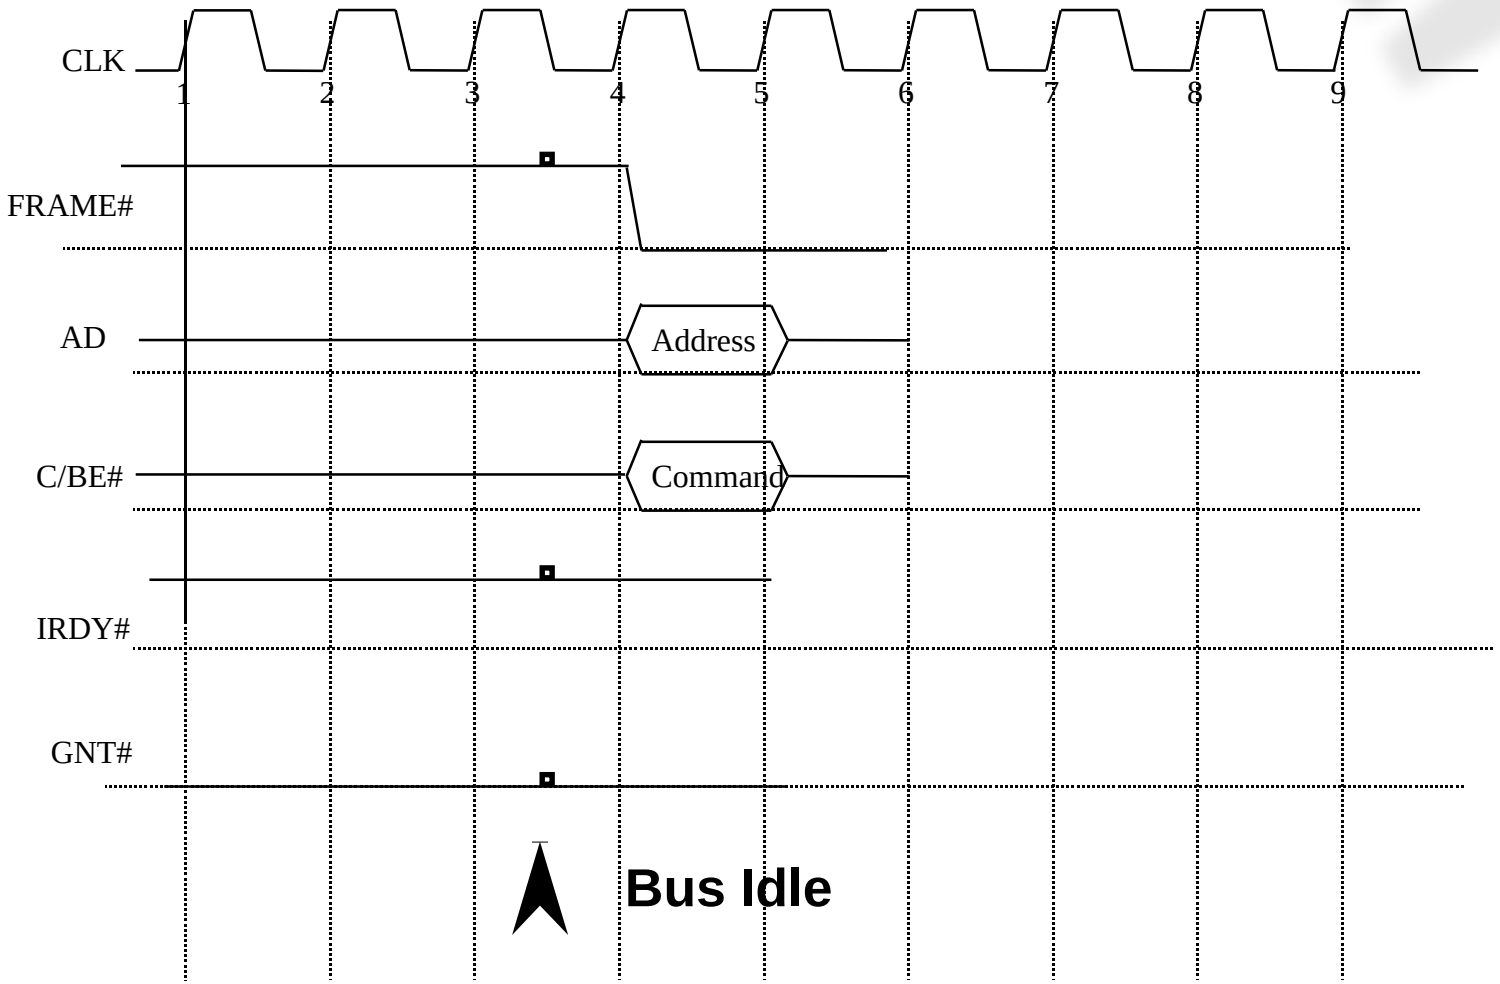

# PCI Command Definition

| C/BE [3::0]# | Command Type |
|---|---|
| 0000 | Interrupt acknowledge |
| 0001 | Special Cycle |
| 0010 | I/O Read |
| 0011 | I/O Write |
| 0110 | Memory Read |
| 0111 | Memory Write |
| 1010 | Configuration Read |
| 1011 | Configuration Write |
| 1100 | Memory Read Multiple |
| 1101 | Dual Address Cycle |
| 1110 | Memory Read Line |
| 1111 | Memory Write and Invalidate |

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Commands

**The PCI bus commands, encoded on the C/BE[3::0]# lines, are summarized below:**

- Interrupt Acknowledge (0000)
  - Performs a read implicitly addressed to the system interrupt controller
  - Address bits are logical don't cares during the address phase
  - Byte enables indicate the size of the vector to be returned
- Special Cycle (0001)
  - Provides a simple message broadcast mechanism on the PCI bus
  - Logical equivalent of a wire to all agents
  - Alternative to physical signals during sideband communications

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Commands (Cont.)

- I/O Read (0010)
  - Reads data from an agent mapped in I/O address space
  - Byte enables indicate the size of the transfer and must be consistent with the byte address
- I/O Write (0011)
  - Writes data to an agent mapped in I/O address space
  - Byte enables indicate the size of the transfer and must be consistent with the byte address
- Reserved (0100, 0101, 1000, 1001)
  - Reserved for future use
- Memory Read (0110)
  - Reads data from an agent mapped in memory address space
  - Target can do anticipatory read if no side effects can be guaranteed (Prefetchable Memory)
  - Target must ensure coherency of data in temporary buffers after this transaction

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Commands (Cont.)

- **Memory Write (0111)**
  - Writes data to an agent mapped in memory address space
  - Target assumes responsibility for coherency of data
- **Configuration Read (1010)**
  - Reads the configuration space of each agent
- **Configuration Write (1011)**
  - Transfers data to the configuration space of each agent
- **Memory Read Multiple (1100)**
  - Fetches a full cache line and starts fetching the next
  - Differs from Memory Read by the intent to fetch more than one cache line before disconnecting

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Commands (Cont.)

- **Dual Address Cycle (1101)**
  - Generates two 32-bit address cycles to produce one 64-bit address
- **Memory Read Line (1110)**
  - Reads up to the cacheline boundary
  - Differs from Memory Read by the intent to complete a full cacheline read
- **Memory Write and Invalidate (1111)**
  - Transfers a complete cacheline to memory, then invalidates the line in cache (Writeback & HITM#)
  - Differs from Memory Write by its guarantee of a minimum transfer of one complete cacheline

QUICKLOGIC
Embedded Standard Products . . . a ger

# Memory Access

- **Reads**
  - **Memory Read** **0110**
    - **burst of less than a cache line**
  - **Memory Read Line** **1110**
    - **burst of cache line**
  - **Memory Read Multiple** **1100**
- **Writes**
  - **Memory Write** **0111**
  - **Memory Write & Invalidate** **1111**

QUICKLOGIC
Embedded Standard Products . . . a ger

# Bus Read

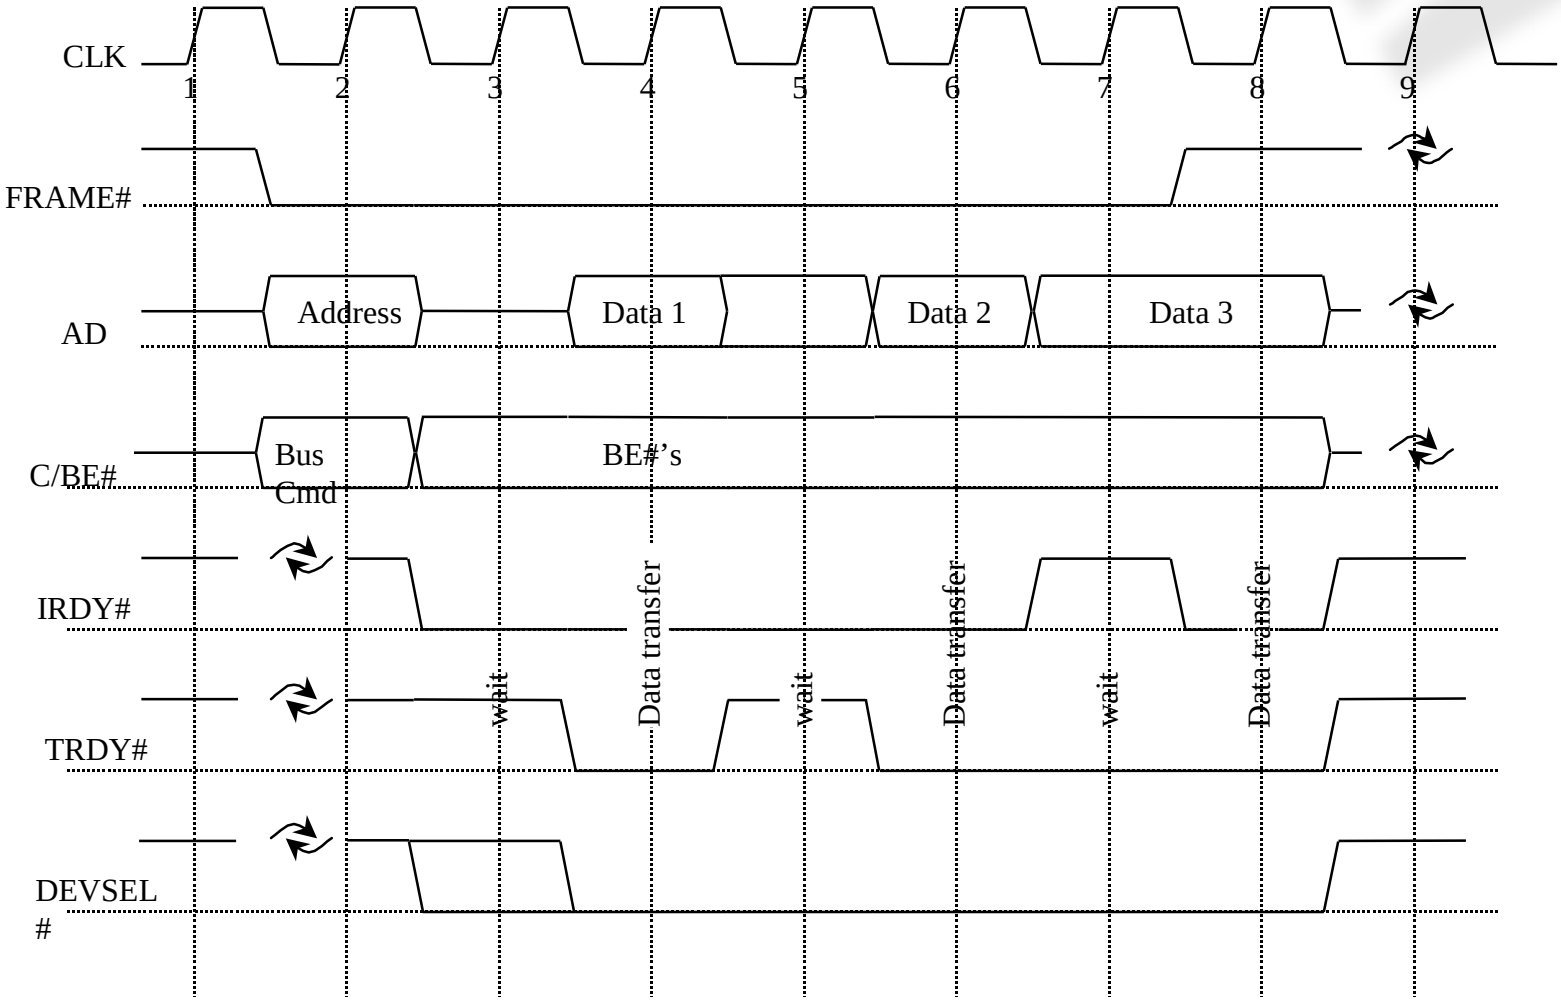
CLK
1
2
3
4
5
6
7
8
9
FRAME#
AD
Address
Data 1
Data 2
Data 3
C/BE#
Bus Cmd
BE#'s
IRDY#
TRDY#
DEVSEL#
wait
Data transfer
wait
Data transfer
wait
Data transfer

# PCI Bus Read

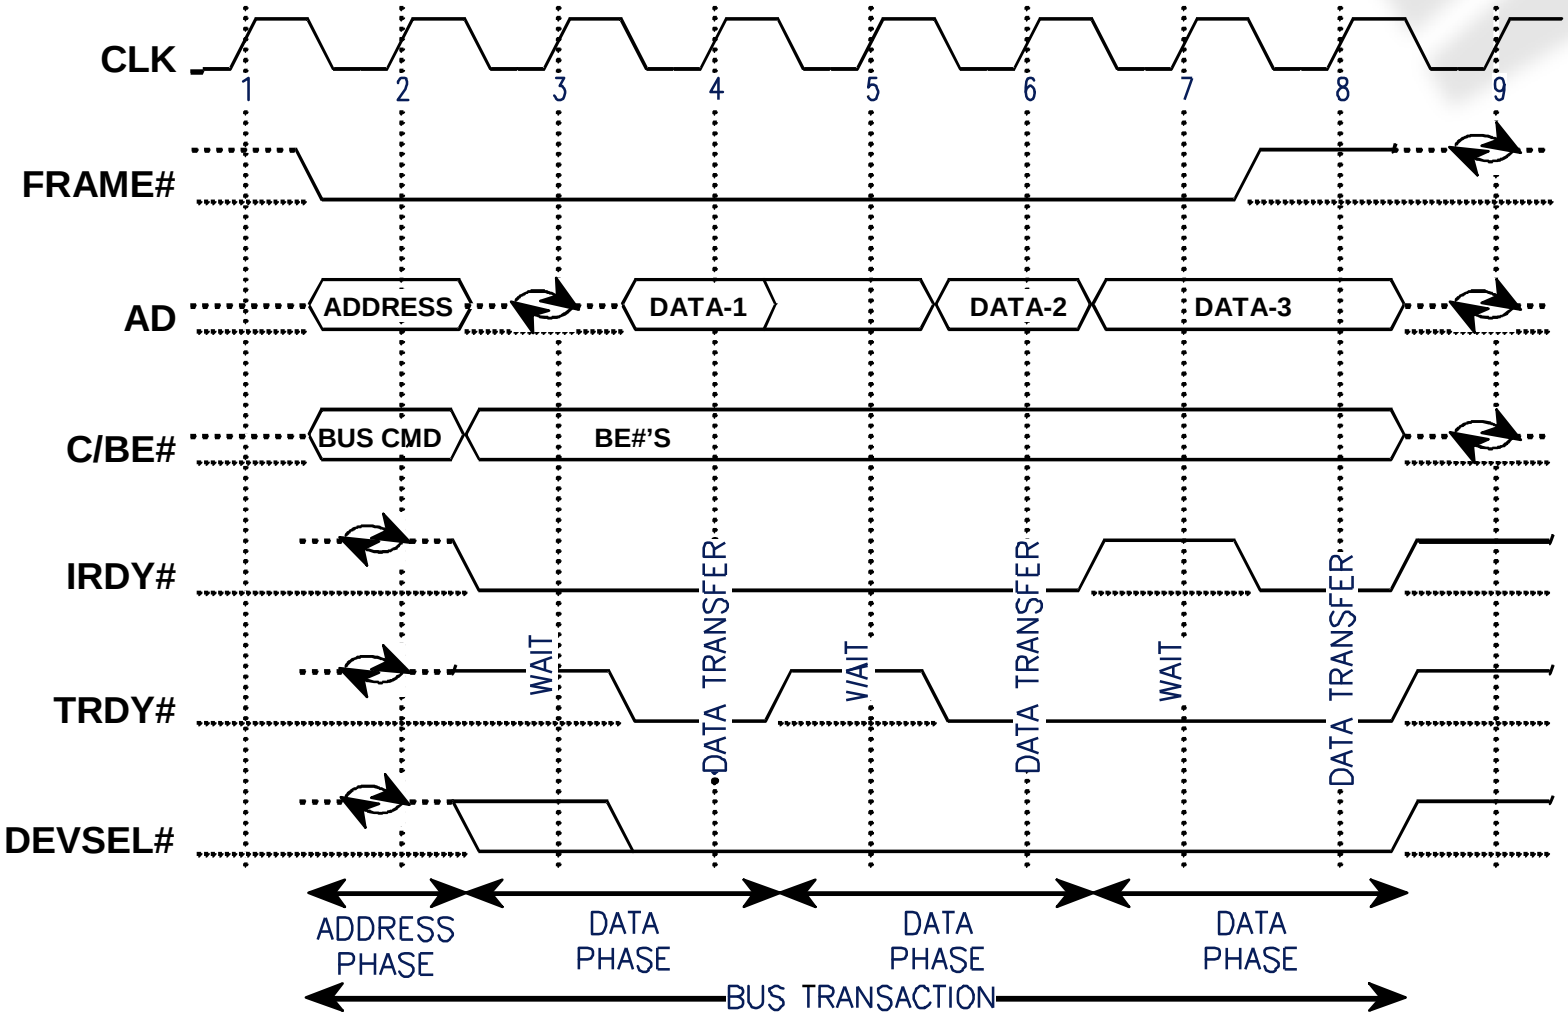
CLK
1
2
3
4
5
6
7
8
9
FRAME#
AD
ADDRESS
DATA-1
DATA-2
DATA-3
C/BE#
BUS CMD
BE#'S
IRDY#
TRDY#
DEVSEL#
WAIT
DATA TRANSFER
WAIT
DATA TRANSFER
WAIT
DATA TRANSFER
ADDRESS PHASE
DATA PHASE
DATA PHASE
DATA PHASE
BUS TRANSACTION

# PCI Bus Write

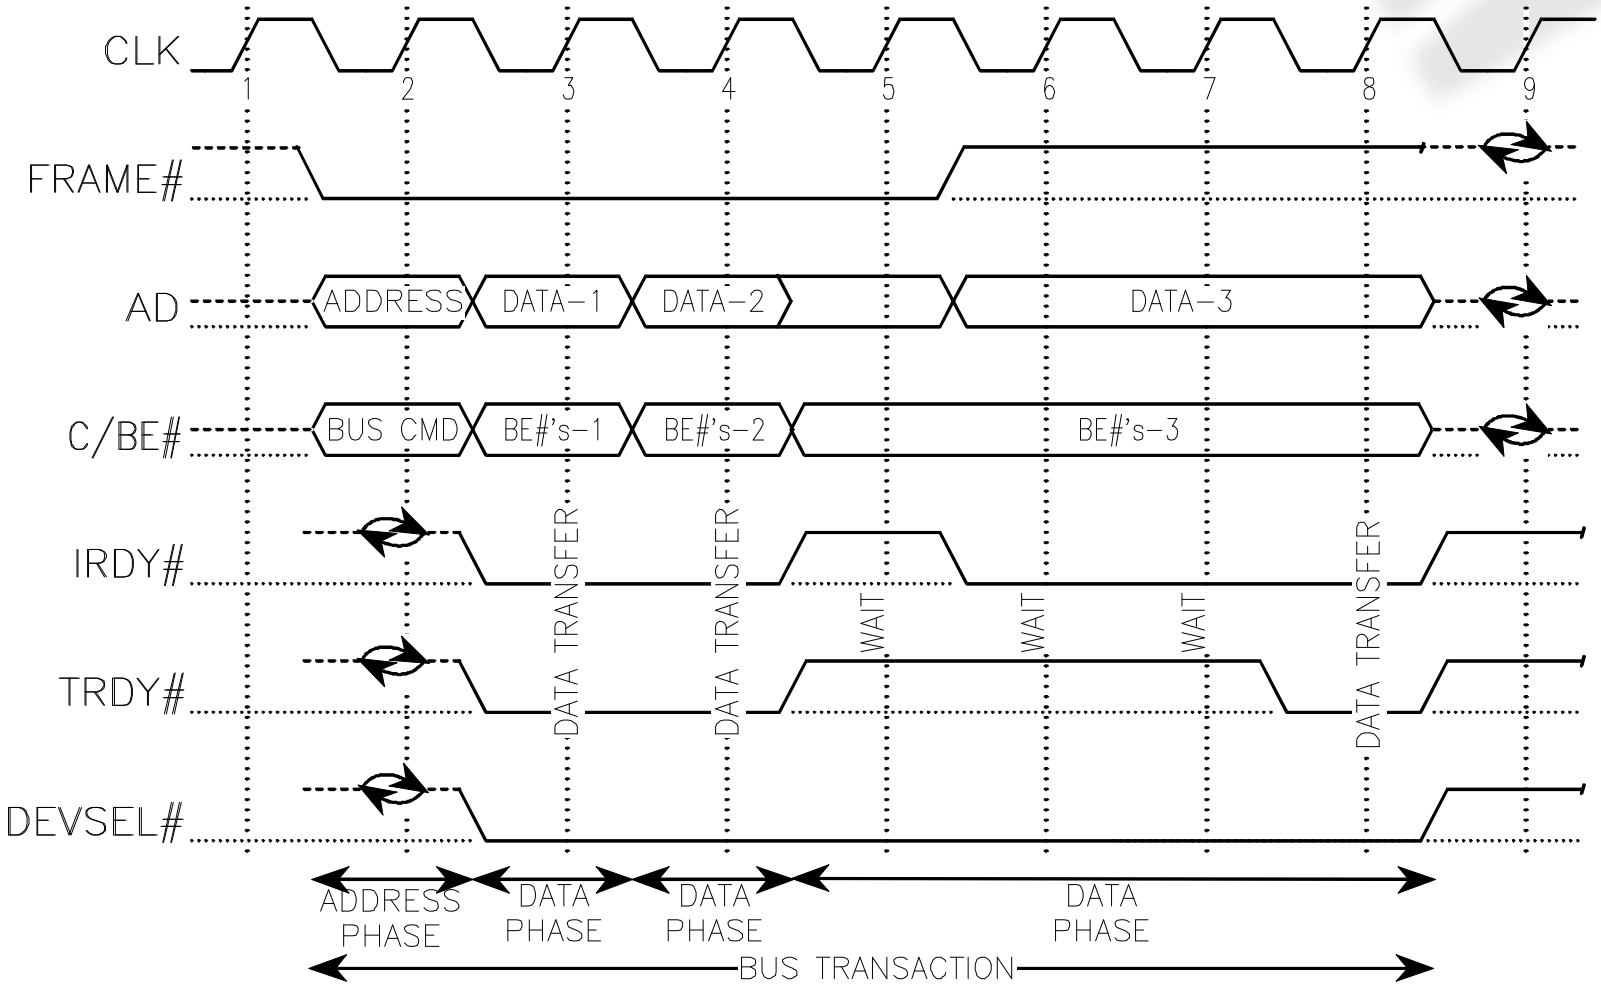

# Premature Termination

- Master Abort
  - No DEVSEL from any target
- Disconnect
  - Target can’t continue burst, stop after current transfer
- Retry
  - Target can’t complete current cycle, retry later
- Abort
  - Target has fatal error, don’t retry

QUICKLOGIC
Embedded Standard Products . . . a ger

# Master Abort

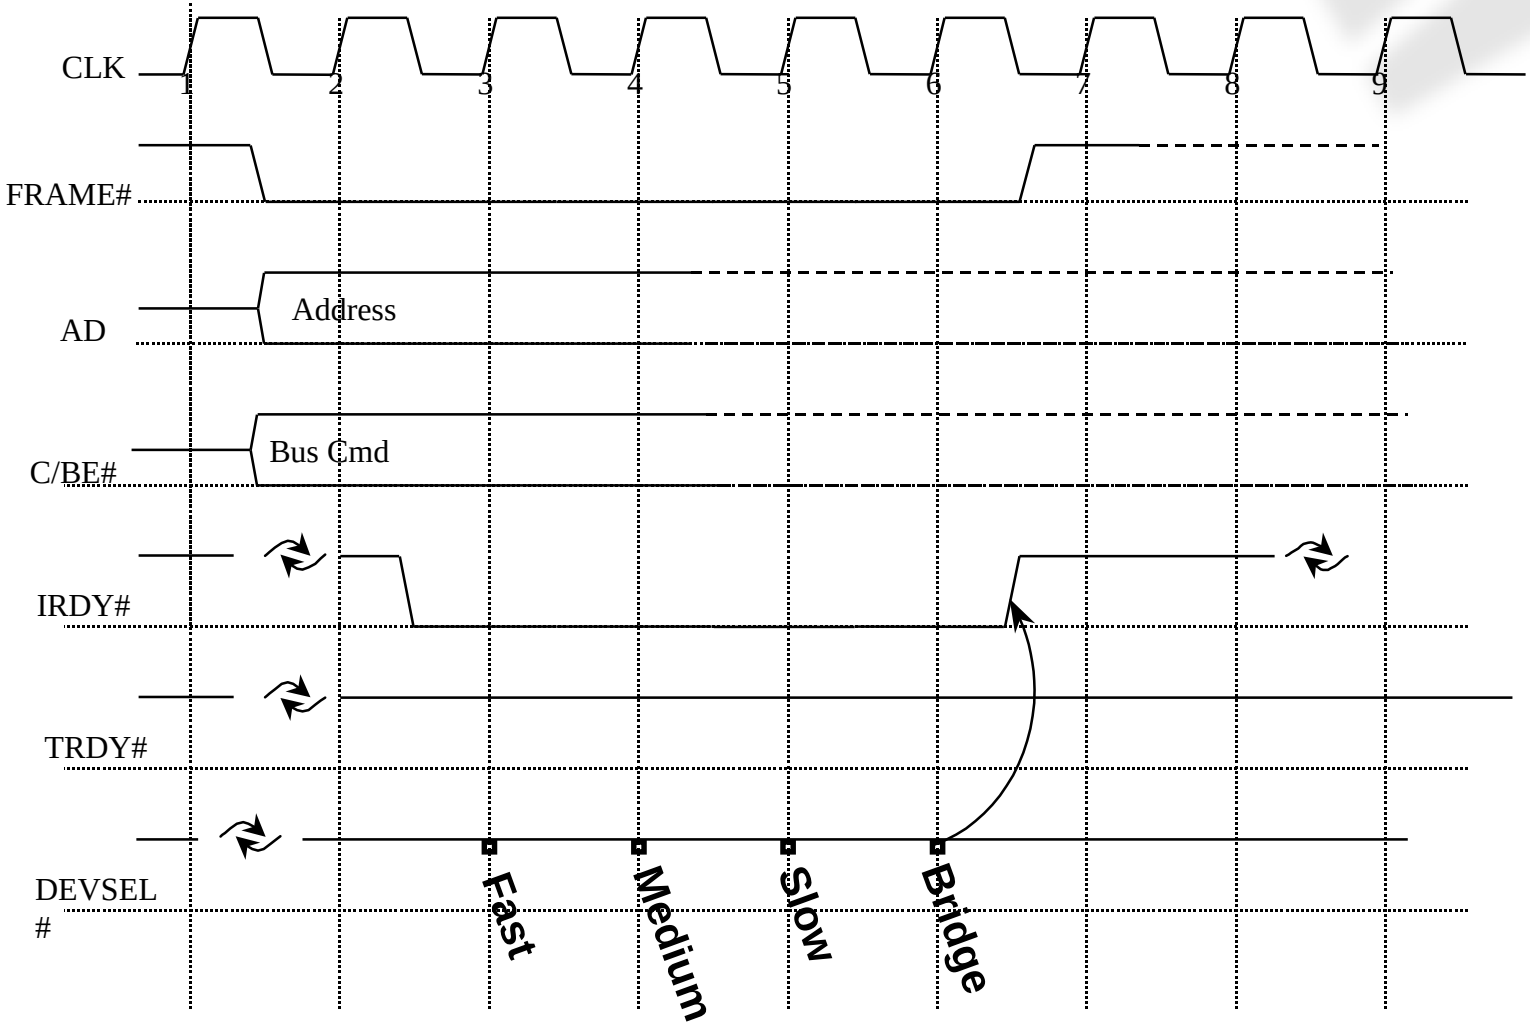

# Disconnect

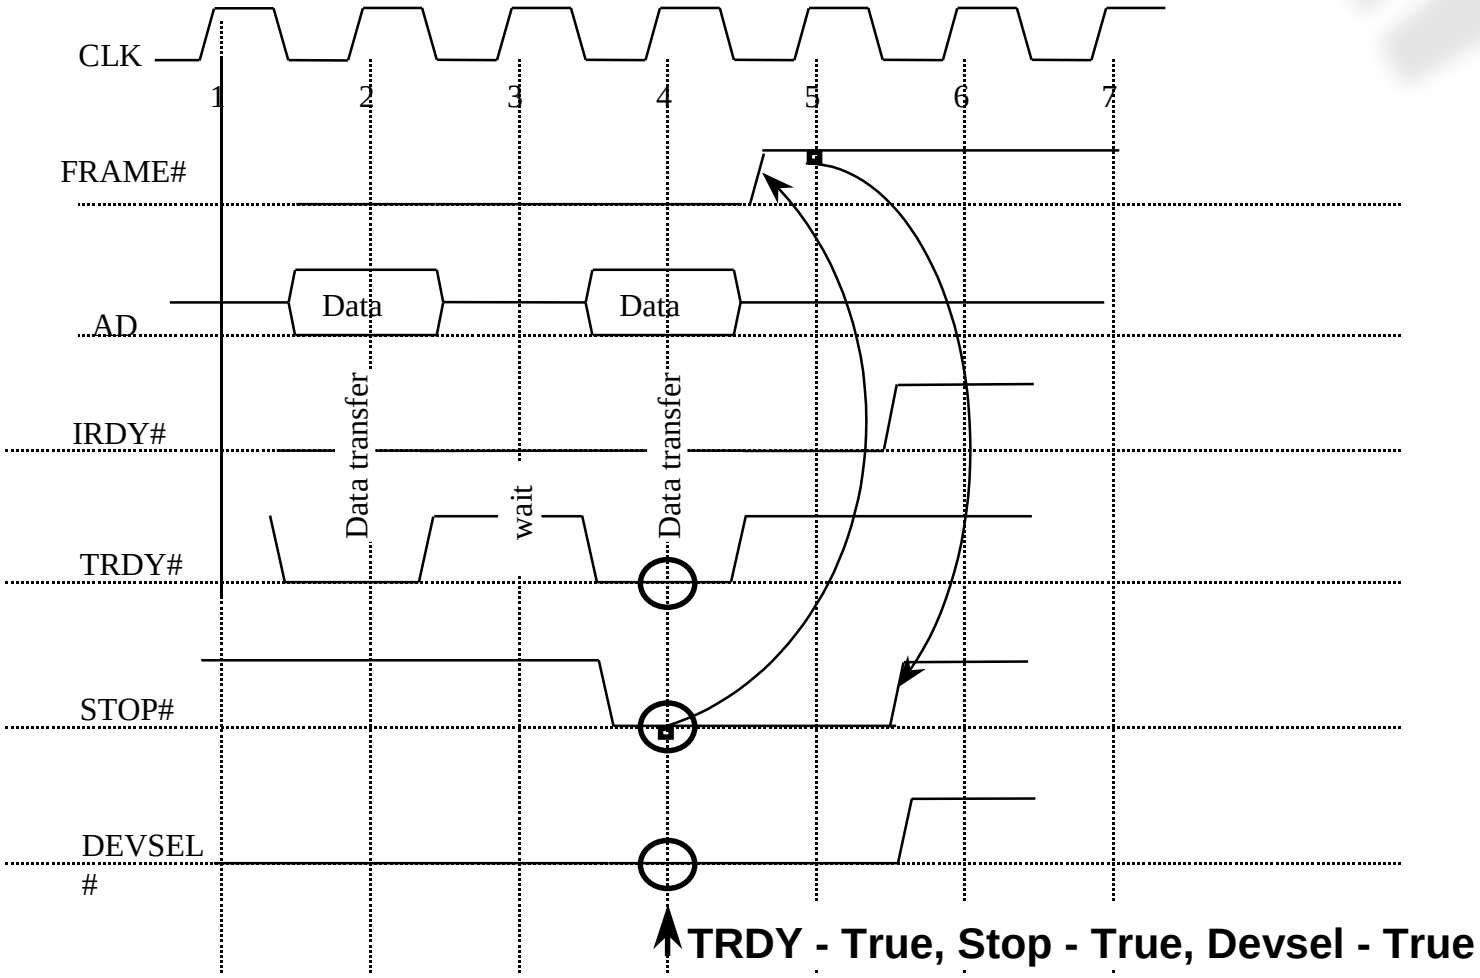
CLK
1
2
3
4
5
6
7
FRAME#
AD
Data
Data
IRDY#
Data transfer
wait
Data transfer
TRDY#
STOP#
DEVSEL
#
TRDY - True, Stop - True, Devsel - True

# Retry

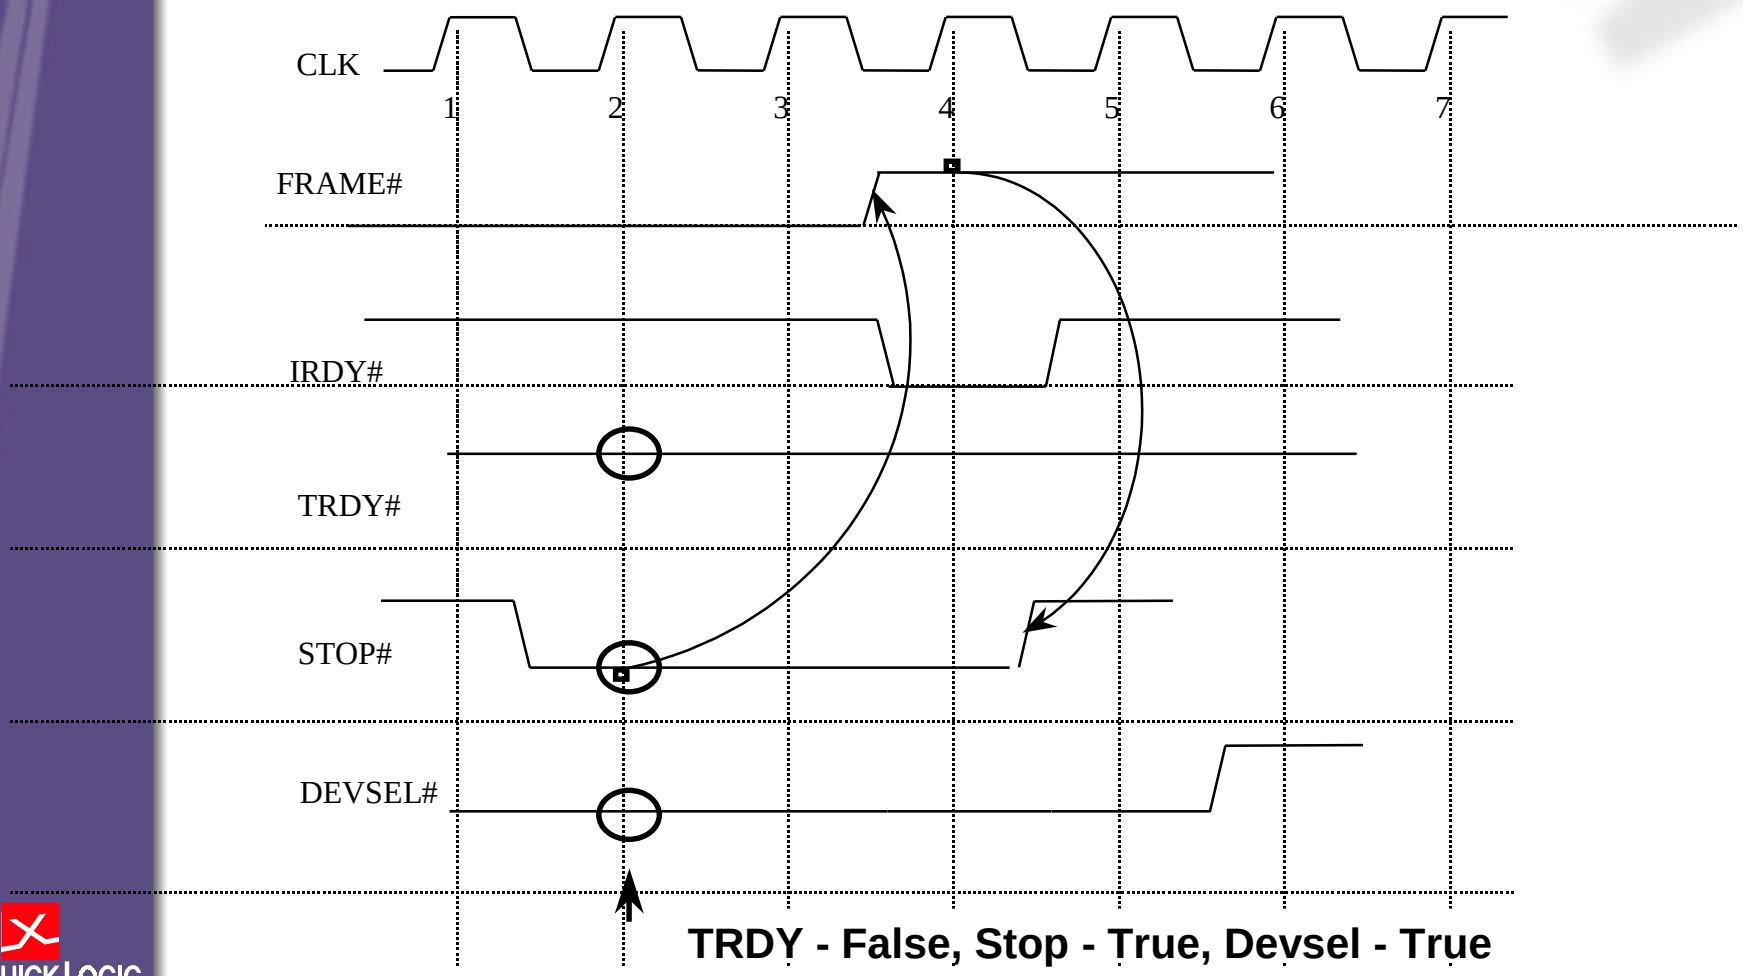


**TRDY - False, Stop - True, Devsel - True**

# Target Abort

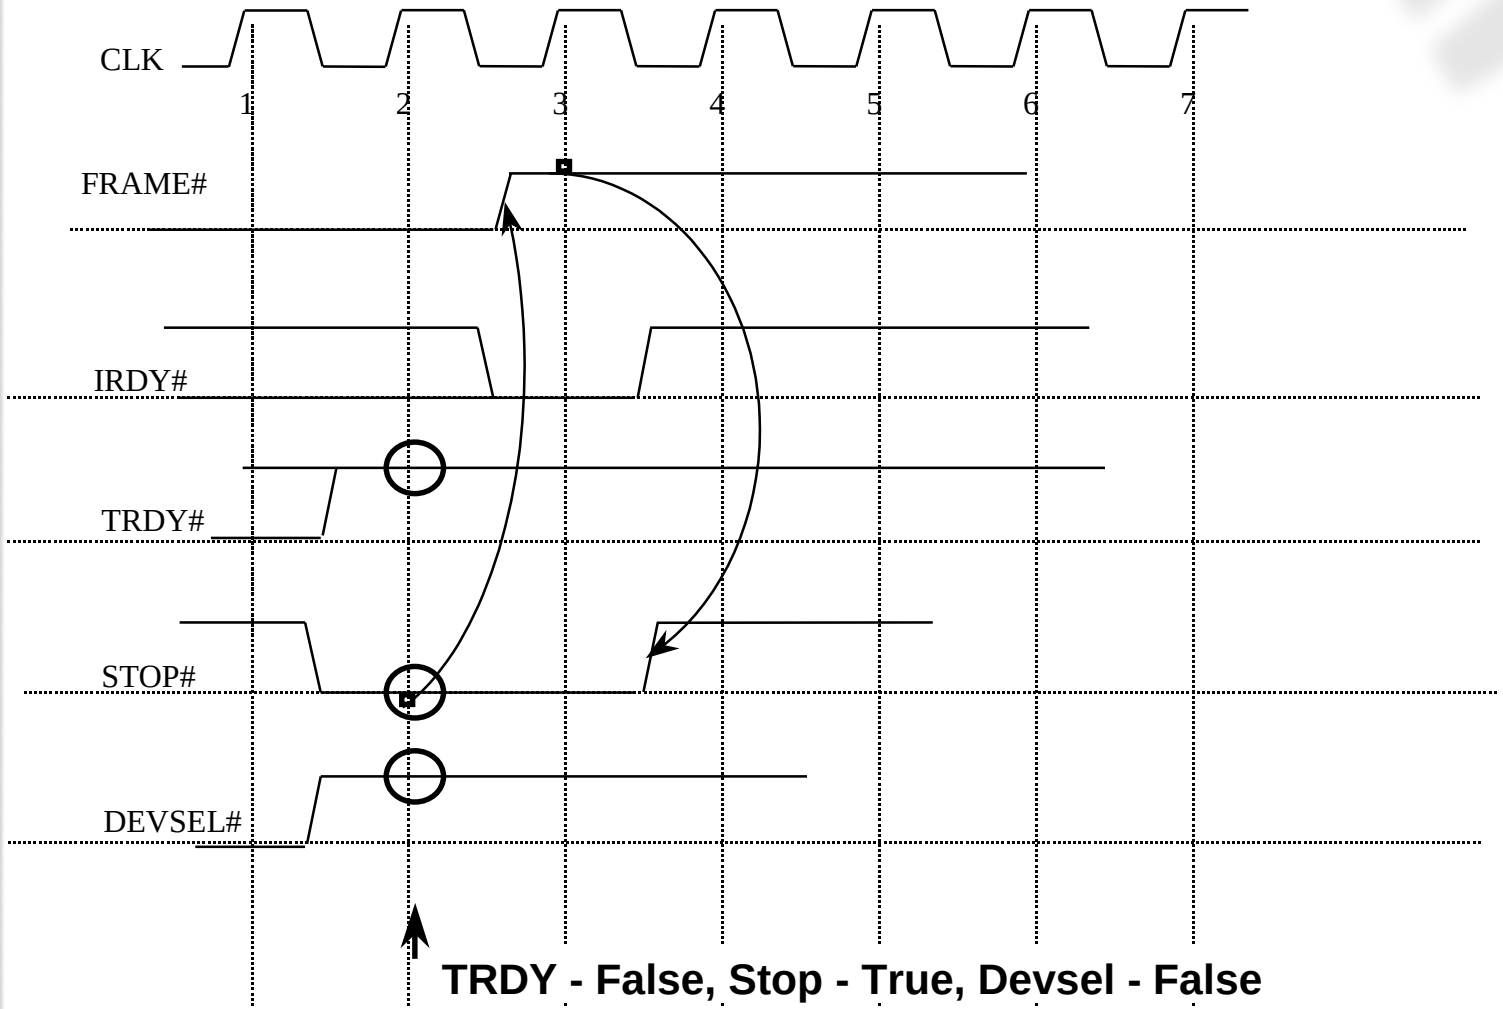


**TRDY - False, Stop - True, Devsel - False**

# PCI Arbitration

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# PCI Bus Arbitration

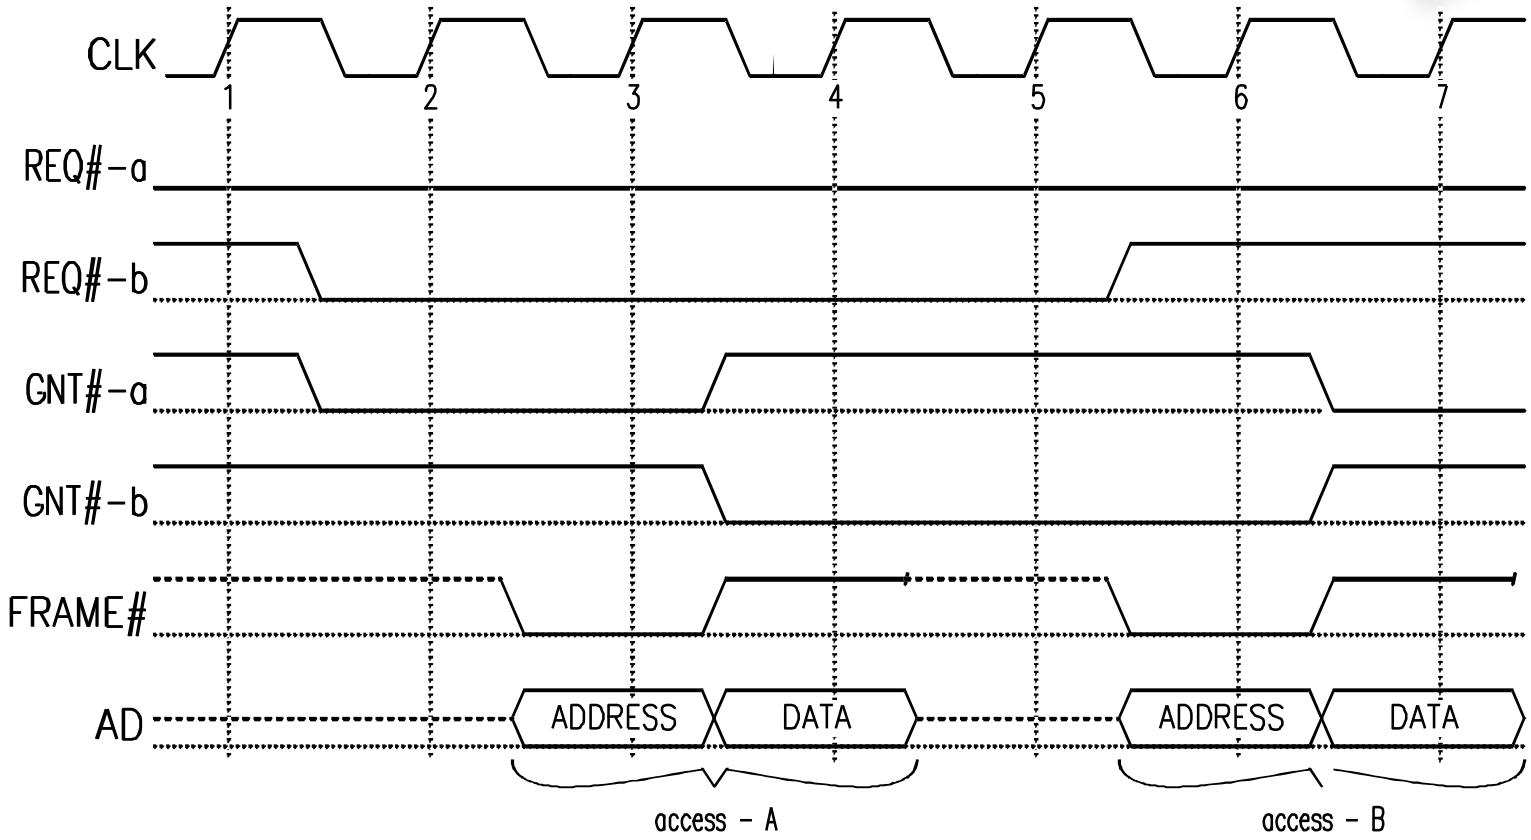

# Arbitration

- Arbitration is access based
  - Master must arbitrate for each bus access
- Central arbitration scheme
  - Each master has a unique request and grant signal
- Arbitration is hidden
  - Occurs during previous bus cycle

QUICKLOGIC
Embedded Standard Products . . . a ger

# Arbitration Scheme

- An arbiter may implement any scheme which is fair
  - Fairness means that, in general, the the arbiter must advance to a new agent when the current master releases the bus
  - Each potential master must be granted  access to the bus independently of other requests

# Bus Parking

- Parking permits the arbiter to select an agent, by asserting its GNT#, when no other agent is using or requesting the bus
- The arbiter determines how this selection is made
  - Fixed, Last Used, …, or None

# Delayed Transactions

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# Delayed Transactions

- Used by two types of devices:
  - I/O Controllers
    - Single delayed transaction in general
  - Bridges
    - Particularly PCI-to-PCI bridges
    - Multiple delayed transactions in general

QUICKLOGIC
Embedded Standard Products . . . a ger

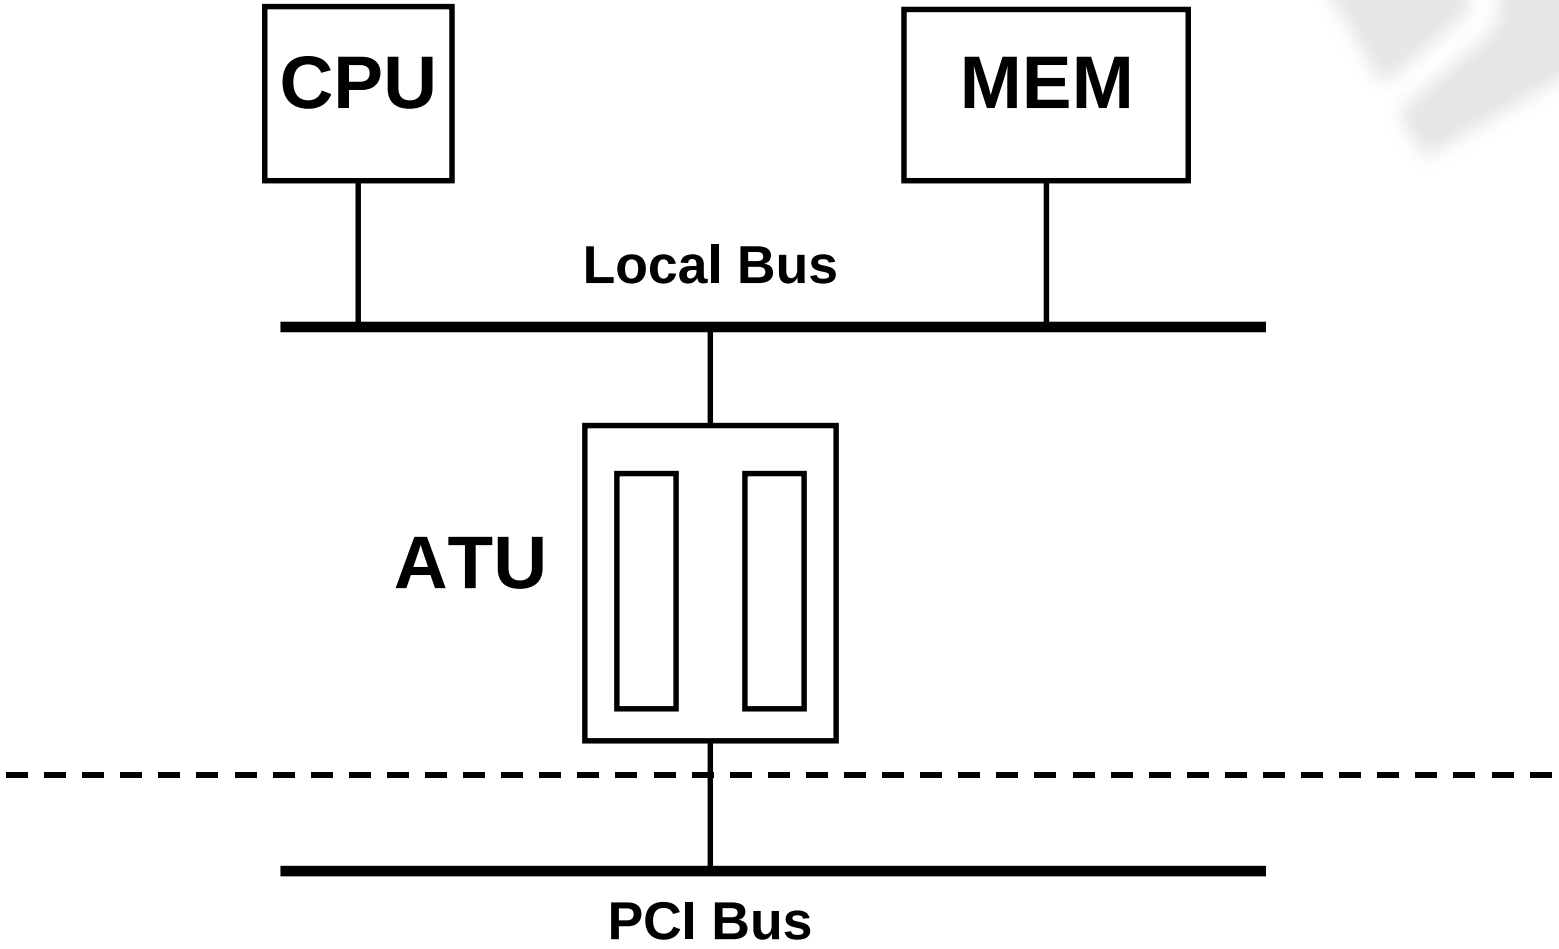
CPU
MEM
Local Bus
ATU
PCI Bus

# Delayed Transactions

- Master issues read
- Slave latches date
- Slave terminates with retry
- Slave transfers data from memory to buffer
- Master issues repeat read
- Slave terminates with TRDY

QUICKLOGIC
Embedded Standard Products . . . a ger

# Latency

QUICKLOGIC

Embedded Standard Products . . . a generation ahead

# Latency

- Arbitration latency
  - Bus acquisition latency
- Master data latency
- Target latency
  - Initial
  - Subsequent

# Bus Access Latency

- Number of clocks from the assertion of REQ# to the completion of the first data transaction
  - Sum of:
    - Arbitration Latency
    - Bus Acquisition Latency
    - Master Latency
    - Target Latency

# Latency

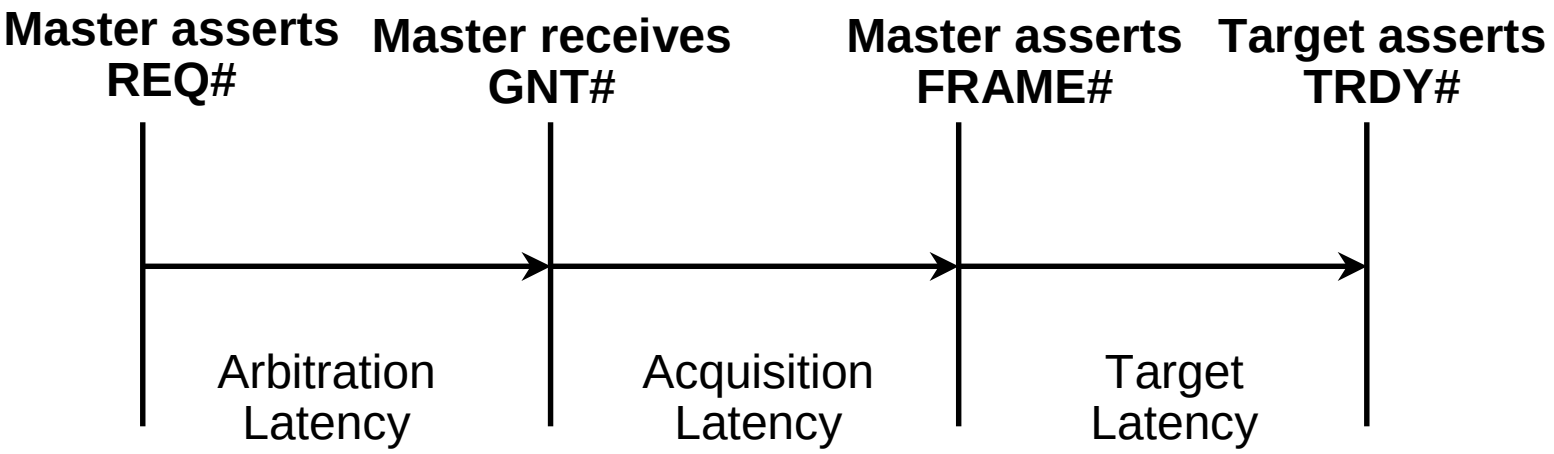
Master asserts REQ#
Master receives GNT#
Master asserts FRAME#
Target asserts TRDY#
Arbitration Latency
Acquisition Latency
Target Latency

# Latency Timer

- Latency Timer (LT) controls trade-off between high throughput and low latency
- Latency Timer counter sets minimum number of clocks before the master must surrender GNT#

# Exclusive Access

## Lock

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# Exclusive Access

- Lock is an exclusive access mechanism which also permits non-exclusive accesses to proceed
- This is referred to as Resource Lock. A hardware lock can be held across several accesses without interfering with other, non-exclusive, real-time transfers
- This mechanism locks only the target of the original locked access. Other transfers, to agents other than the locked device, may share the bus
- Exclusive access is provided through the use of the LOCK# signal

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Hardware

# PCI Card Connectors

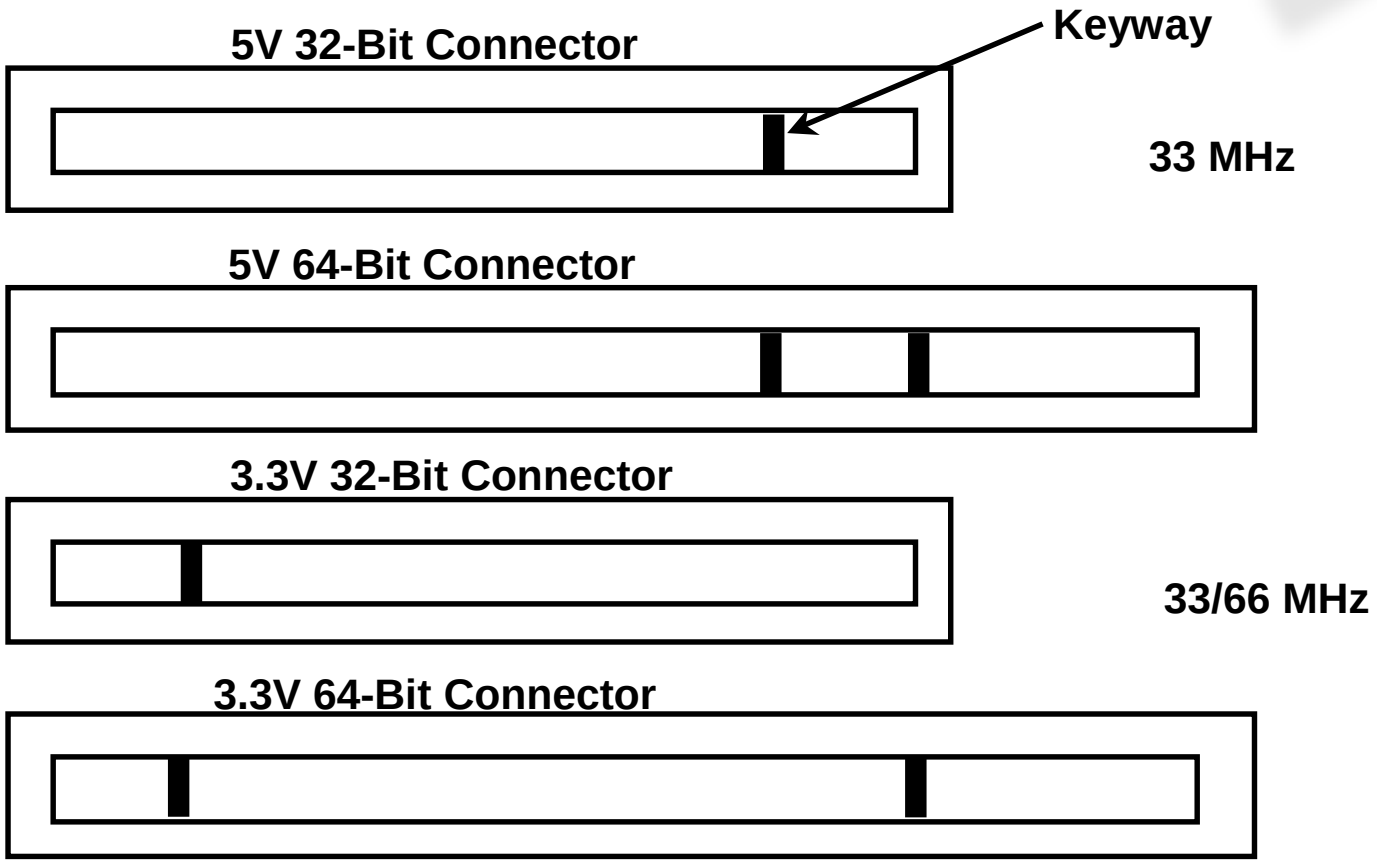

# 5 V To 3.3 V Migration Path

- ***5 Volt Card***
- ***Dual-Voltage Card***
- ***3 Volt Card***

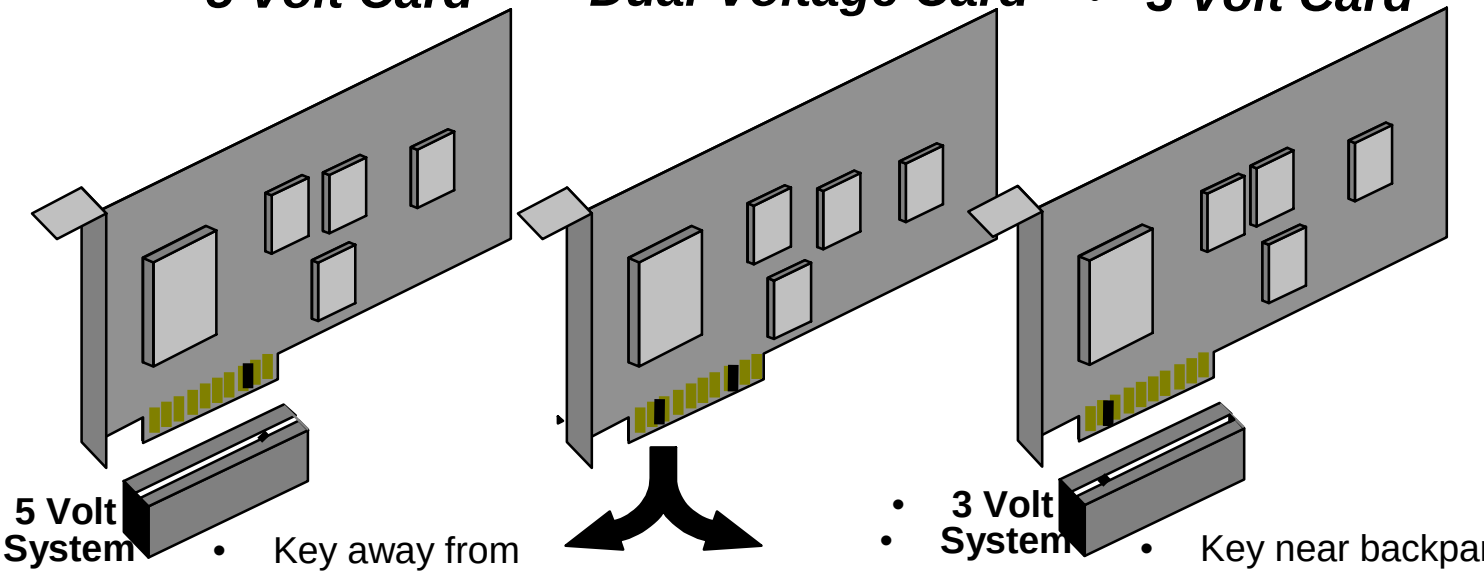

- **5 Volt System**
- Key away from backpanel
- **3 Volt System**
- Key near backpanel

# PCI Component Pin Out

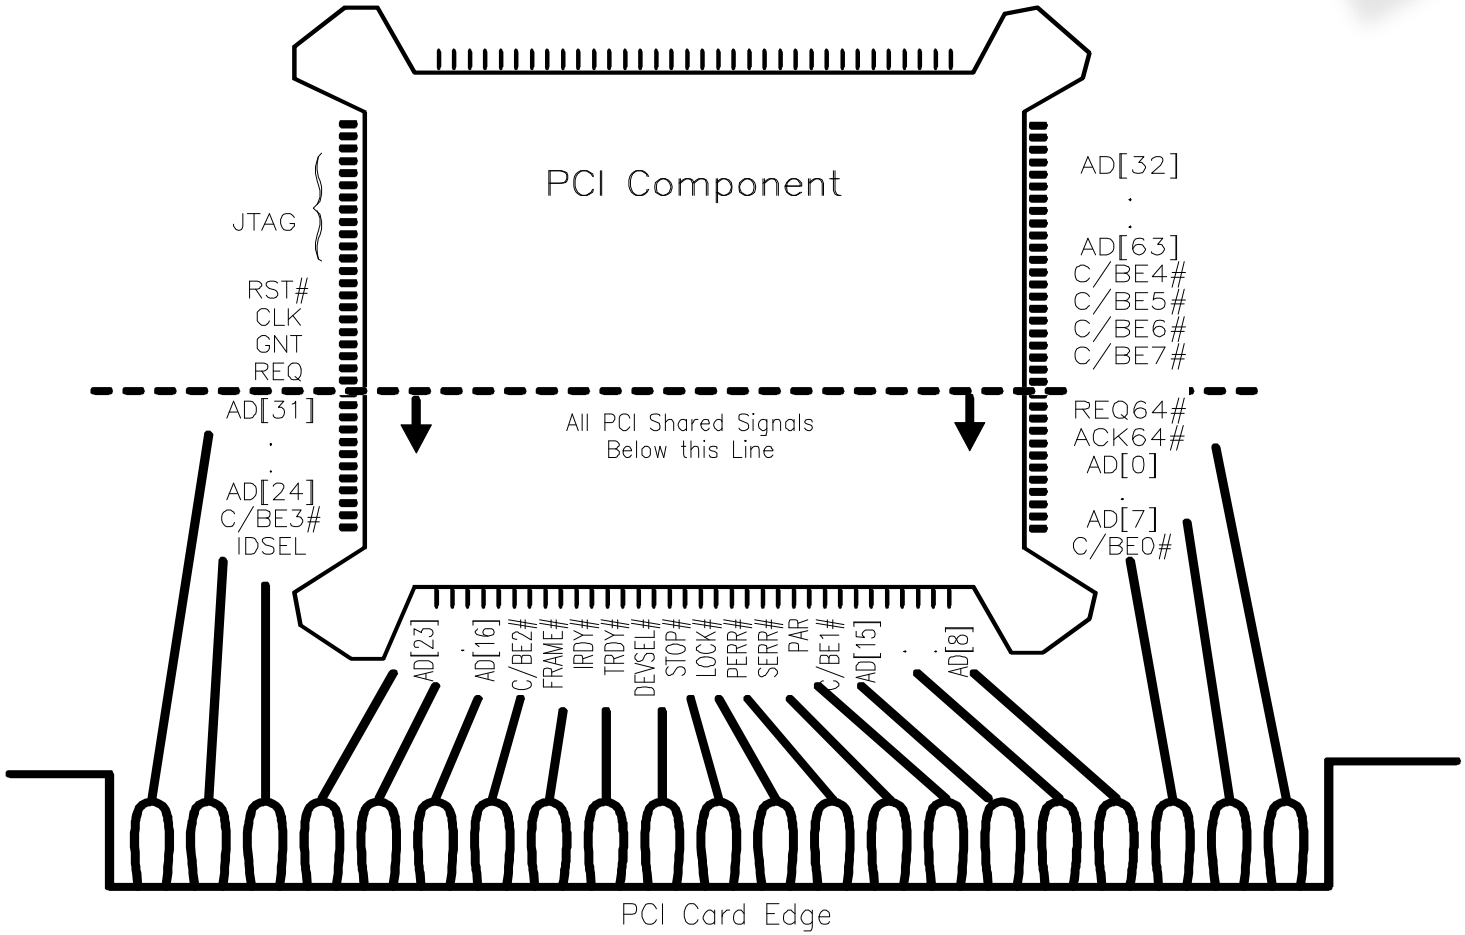

# Trace Length

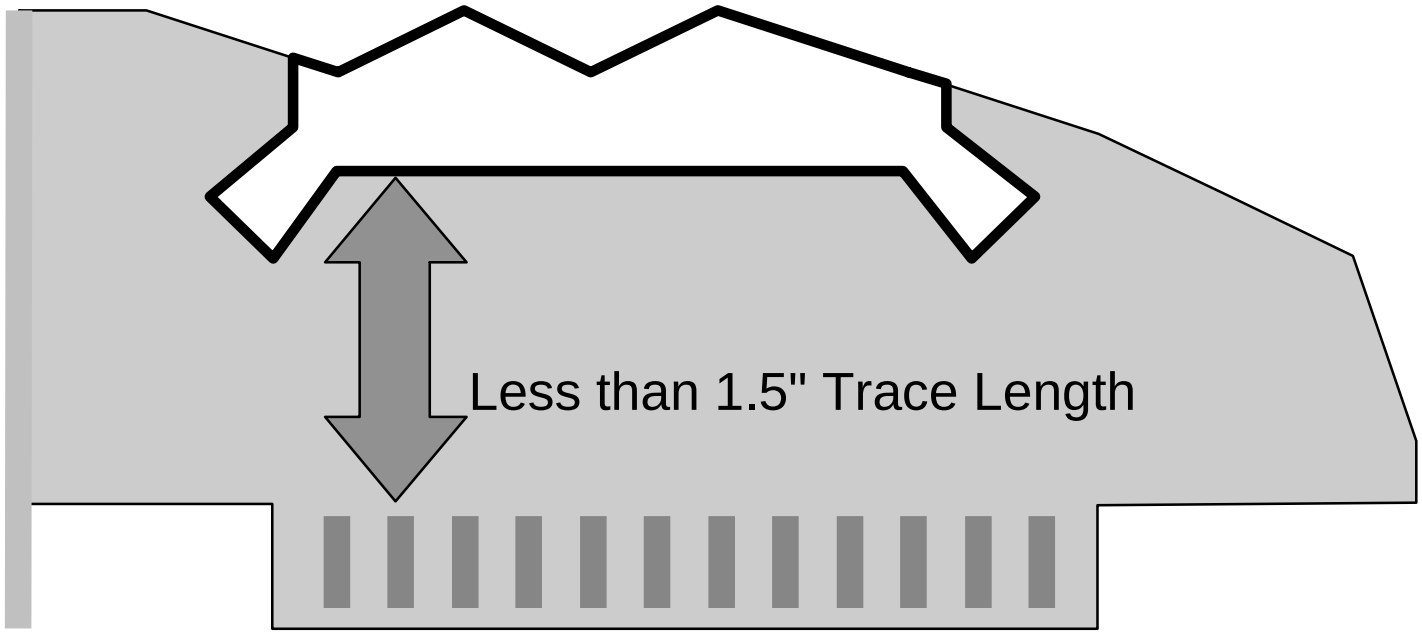

# PCI Component Bus Loading

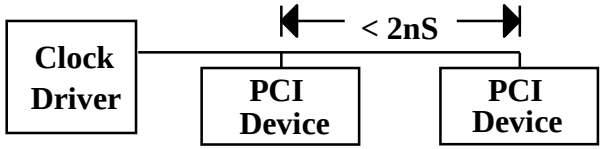


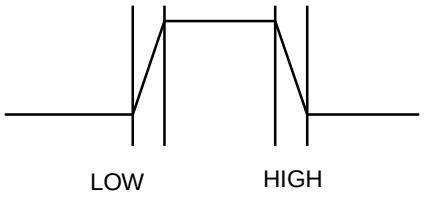


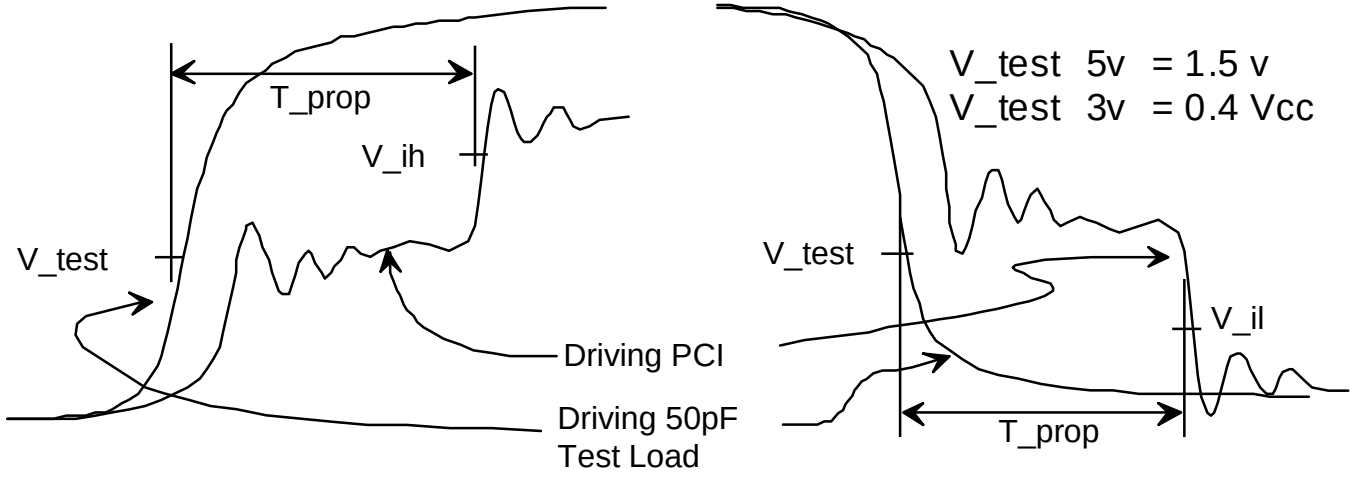

# PCI Load Definition

- Tprop = 10ns @33MHz, 5ns @ 66MHz
- <10 PF per PCI line
- <20 NH per PCI line
- 60 - 100 Ohms Z
- 150 to 190 ps/inch signal velocity (50%C ± 5%)
- Shared signals are limited to 1.5 inches
- Unshared signals are limited to 2 inches
- Clock must be 2.5 ± .1 inch, only one load
- Max Clock skew 2ns @ 33MHz, 1ns @ 66 MHz
- No maximum number of loads is given in the specification. Pullup resistor calculations assume 10 to be typical, and 16 to be worst case loading

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Load Restrictions

It is specifically a violation of the PCI specification to:

- Attach an expansion ROM directly, or via bus transceivers, on any PCI pins
  - ROM code may not execute in-place
- Attach two or more PCI devices on an expansion board except behind a PCI component
- Attach any logic, other than a single PCI device, which snoops PCI pins
- Use PCI component sets which place more than one load on any PCI pin,
  - Separate address and data path components
- Use a PCI component which has more than 10pF per pin
- Attach any pull-up resistors, or other discrete devices, to PCI signals except behind PCI component

QUICKLOGIC
Embedded Standard Products . . . a ger

# 66MHz Operation

- PCI operates at 66 MHz only in the 3.3V environment
- Agents indicate 66MHz capability both electrically and programmatically
  - M66EN pin
  - 66M bit in status register
- Any 33MHz board on the bus grounds the M66EN pin and disables 66 MHz operation
  - Pin B49 grounded indicates 33MHz

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI BIOS

## 32-Bit Operating Systems

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# PCI BIOS

- **Only BIOS is permitted to access PCI:**
  - Configuration Registers
  - Interrupt Routing Logic
  - Special Cycle generation logic

# PCI BIOS Functions

- Determine configuration mechanism
- Detect all PCI buses in the system
- Search for PCI devices
- Read/Write configuration registers
- Assign Interrupts to PCI devices
- Determine Special Cycle abilities
- Permit Special Cycle generation

QUICKLOGIC
Embedded Standard Products . . . a ger

# PCI Configuration

QUICKLOGIC
Embedded Standard Products . . . a generation ahead

# System Initialization

- Configuration allows software (BIOS) to initialize the system
- Each device has configuration registers
- At power up software scans bus(es)
- Software analyses system requirements
- Configuration registers are set to configure individual devices

QUICKLOGIC
Embedded Standard Products . . . a ger

# Configuration Types

- Specific bus commands
  - configuration read (C/BE# = 1010)
  - configuration write (C/BE# = 1011)
- Type 0
  - local PCI bus
    - IDSEL line indicates device
    - address field indicates register
- Type 1
  - remote PCI bus (through bridge)
    - address field indicates bus, device and register

# Configuration Space Header

| | | | | |
|---|---|---|---|---|
| Device ID | | Vendor ID | | 00 |
| Status | | Command | | 04 |
| Class Code | | | Rev | 08 |
| BIST | Header Type | Latency Timer | Cache Line Size | 0C |
| Base Address 0 | | | | 10 |
| Base Address 1 | | | | 14 |
| Base Address 2 | | | | 18 |
| Base Address 3 | | | | 1C |
| Base Address 4 | | | | 20 |
| Base Address 5 | | | | 24 |
| Cardbus CIS Pointer | | | | 28 |
| Subsystem ID | | Subsystem Vendor ID | | 2C |
| Expansion ROM Base Address | | | | 30 |
| Reserved | | | | 34 |
| Reserved | | | | 38 |
| Max_Lat | Min_Gnt | Interrupt Pin | Interrupt Line | 3C |

# Command Register

| Bit | Meaning |
|---|---|
| 0 | I/O Space Enable |
| 1 | Memory Space Enable |
| 2 | Bus Master Enable |
| 3 | Special Cycle Enable |
| 4 | Memory Write and Invalidate Enable |
| 5 | VGA Palette Snoop Enable |
| 6 | Parity Error Enable |
| 7 | Wait Cycle Control |
| 8 | SERR# Enable |
| 9 | Fast Back to Back Enable |
| 10:15 | Reserved |

QUICKLOGIC
Embedded Standard Products . . . a ger

# IDSEL Line Addressing

- Addressing the Configuration Space of a device is done using the IDSEL signal as a chip select, and setting AD[1::0] to 00 to indicate a Type 0 configuration transaction
- The method used to generate IDSEL is system specific, however, if no other mapping is required the following may be used
- The IDSEL of Device 0 is connected to AD[16], Device 1 is connected to AD[17], etc. through Device 16 connected to AD[31]

QUICKLOGIC
Embedded Standard Products . . . a ger

# Interrupt Line

- 8-Bit register which contains interrupt line routing
- Value in this register is the interrupt number (IRQ) to which the device's interrupt pin is connected
- If, during initialization, the device requests an IRQ, one is assigned.
  - The IRQ number is placed in this register by configuration software at initialization time

QUICKLOGIC
Embedded Standard Products . . . a ger

# Interrupt Pin

- The interrupt pin register indicates which of the four interrupts, INTA#-INTD#, the device or function uses
- A value of 1 corresponds to INTA#, 2 to INTB# etc
- Devices which do not use interrupts place 0 in this register
- This register is read only
- Each function within a device may use one interrupt
  - A single function device must use INTA#
  - If a multifunction device requires two interrupts they must be INTA# and INTB# etc.
- Interrupts may be shared (chained)

# Interrupt Routing

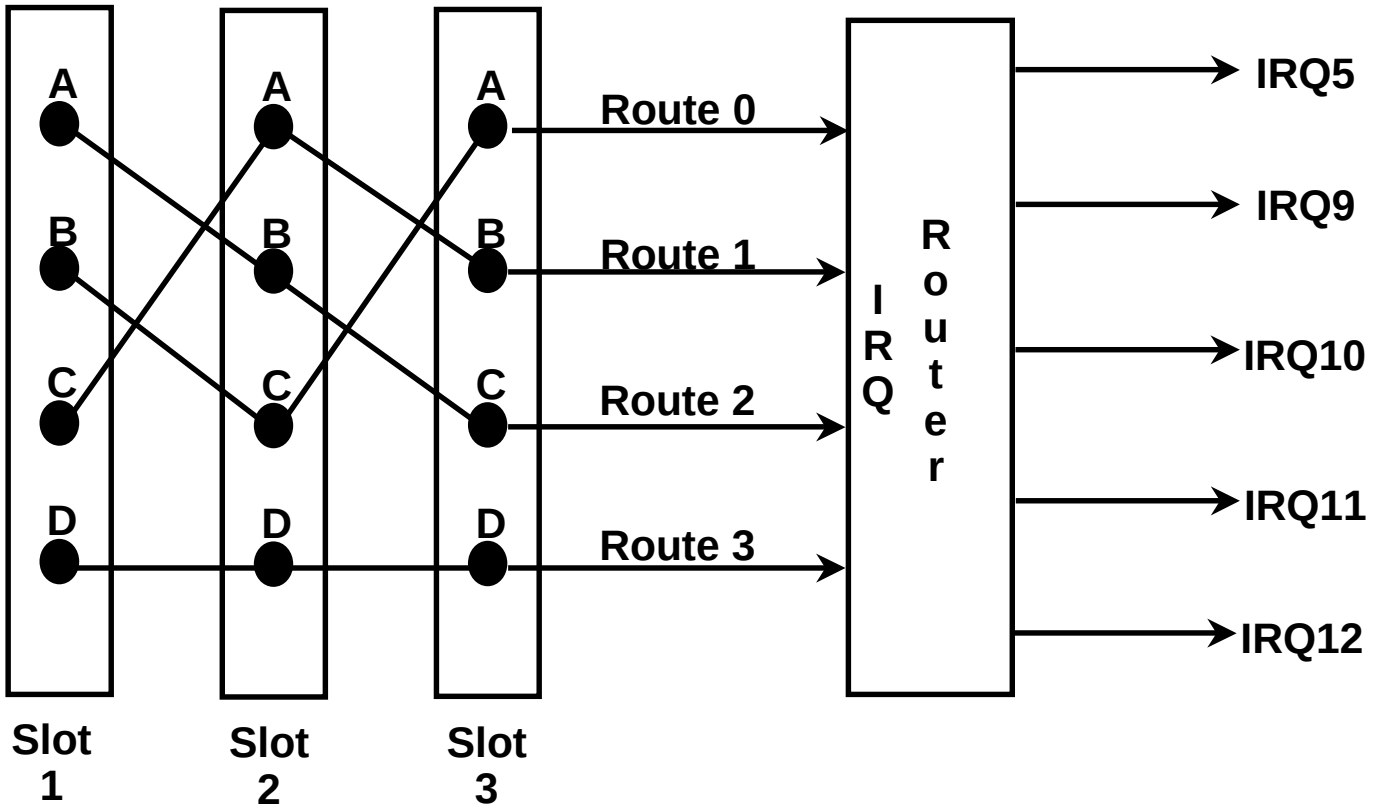

# PCI Bus Interrupts

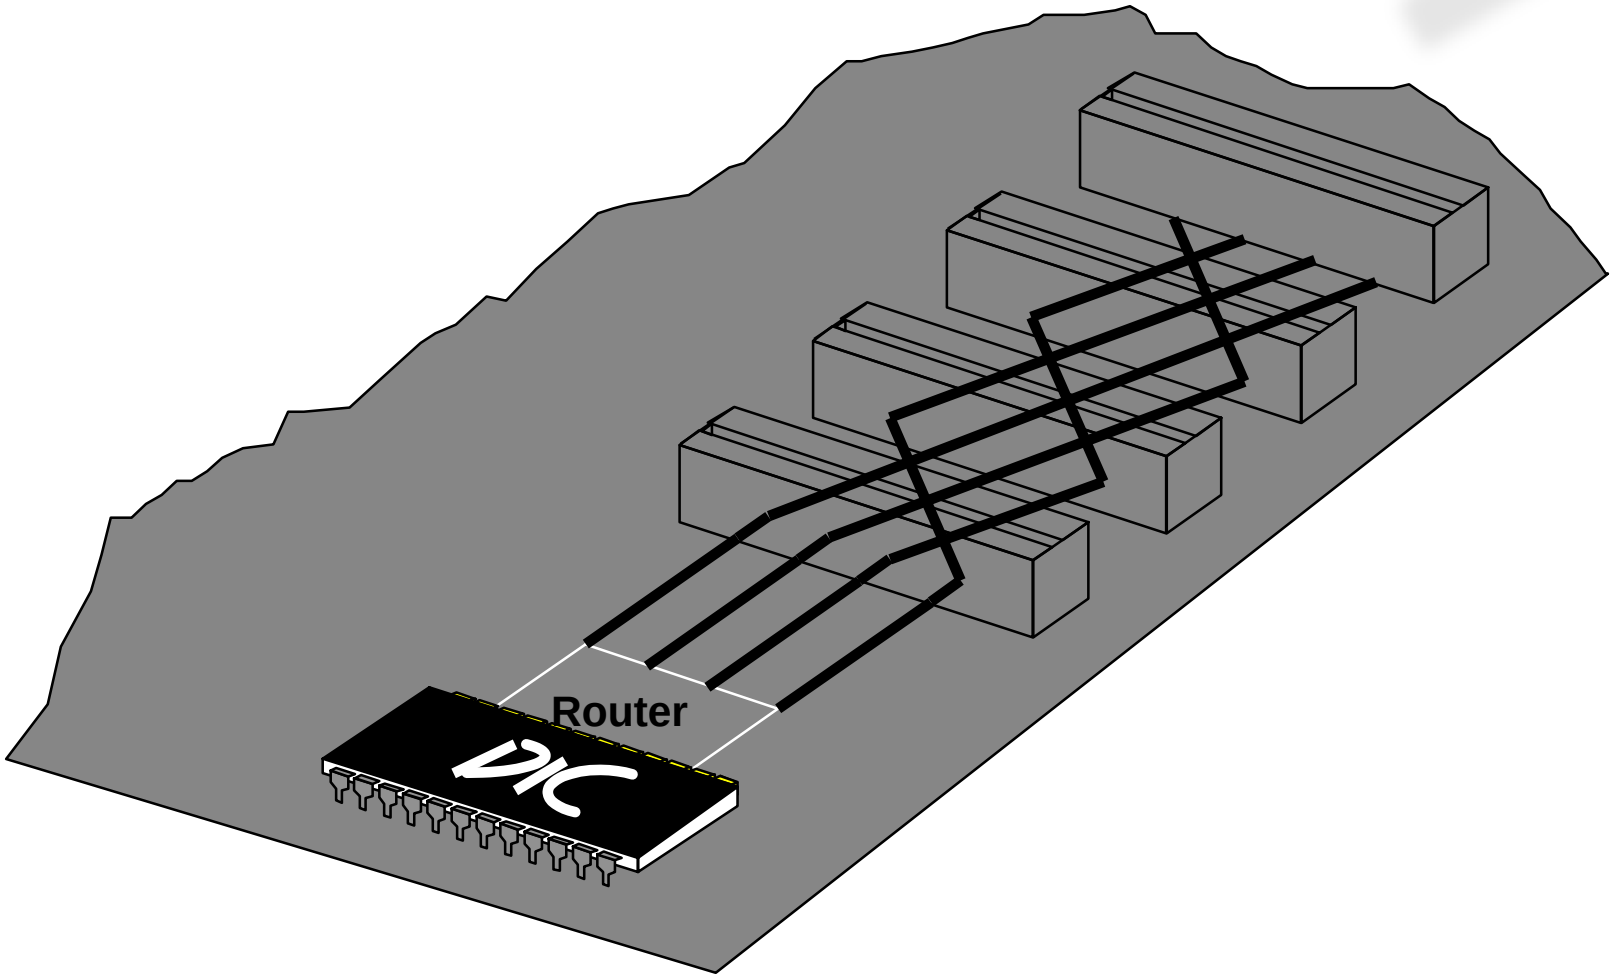

# MIN_GNT and MAX_LAT

- Minimum Grant and Maximum Latency
- MIN_GNT defines the burst-period length required by the device
- MAX_LAT specifies how often the device must gain access to the PCI bus
  - Maximum time between accesses
- Both registers are read only and have a granularity of 250 nanoseconds
- Values of 0 in these registers indicates no major latency requirements for the device

QUICKLOGIC
Embedded Standard Products . . . a ger

# Documentation

- PCI Local bus Specification Rev 2.1
  - PCI Special Interest Group
  - http://www.pcisig.com/
- PCI System Architecture, 3rd Edition
  - Mindshare
  - ISBN# 0-201-40993-3
- PCI Hardware & Software Architecture and Design
  - Edward Solari & George Willse

QUICKLOGIC
Embedded Standard Products . . . a ger