3. Computer Buses

- Introduction
- Electrical Considerations
- Data Transfer Synchronization
- Parallel and Serial Buses
- Bus Arbitration
- PCI Bus
- PCI Express Bus
- Other Serial Buses
- VME Bus

Parallel and Serial Buses (1)

- Parallel buses
  - Use multiple lines to transmit data words
  - Examples: PCI, VME
  - Technological problems make difficult to increase their clock frequency
    - Timing skew: difference between propagation delays of signals on various lines
- Serial buses
  - Use a single line to transmit data bit by bit
  - Examples: PCI Express, I^2C, SPI, USB

Parallel and Serial Buses (2)

- Clocking information may be embedded within the serial data stream
- Advantages of serial buses:
  - Smaller size of connectors and cables
  - Reduced electrical interference
  - Simpler synchronization
  - Longer interconnection distance
  - Higher reliability
  - Lower cost of the interface

3. Computer Buses

- Introduction
- Electrical Considerations
- Data Transfer Synchronization
- Parallel and Serial Buses
- Bus Arbitration
- PCI Bus
- PCI Express Bus
- Other Serial Buses
- VME Bus

Bus Arbitration

- Bus Arbitration
  - Centralized Bus Arbitration
  - Decentralized Bus Arbitration

Bus Arbitration

- Function: to determine the module that will become *master* in case of simultaneous requests
- Arbitration methods
  - Centralized: bus allocation is performed by a bus arbiter
  - Decentralized (distributed): there is no bus arbiter

Bus Arbitration

- Bus Arbitration
  - Centralized Bus Arbitration
  - Decentralized Bus Arbitration

Centralized Bus Arbitration (1)

- Methods for centralized bus arbitration:
  - Daisy-chaining of devices
  - Independent requesting
  - Polling
- Centralized arbitration using daisy-chaining of devices
  - A single bus request line, *BUSREQ* (*Bus Request*) → wired OR
  - A bus grant line *BUSGNT* (*Bus Grant*)

Centralized Bus Arbitration (2)

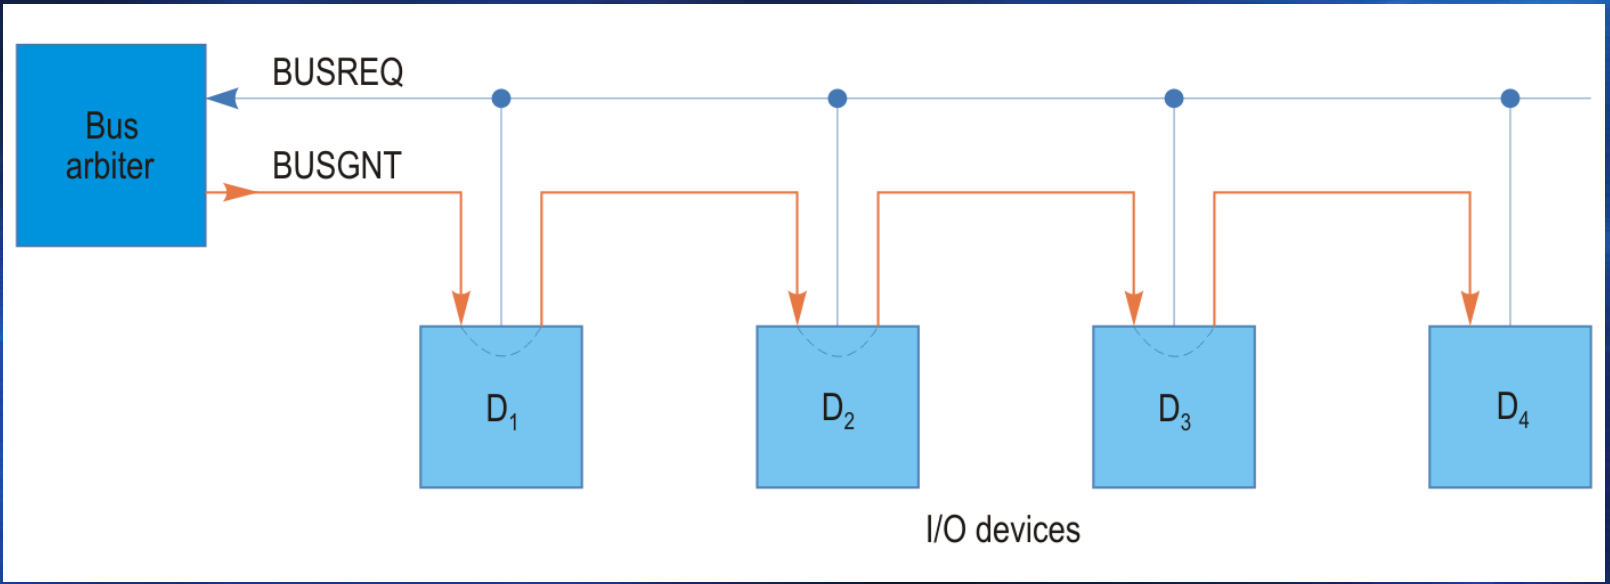

- The device physically closest to the arbiter detects the signal on the *BUSGNT* line

Centralized Bus Arbitration (3)

- Only two control lines are required for bus arbitration
- Device priority is fixed ⟶ given by the chaining order on the *BUSGNT* line
- To modify the default priorities, buses may have multiple priority levels
  - For each priority level, there is a bus request line and a bus grant line

Centralized Bus Arbitration (4)

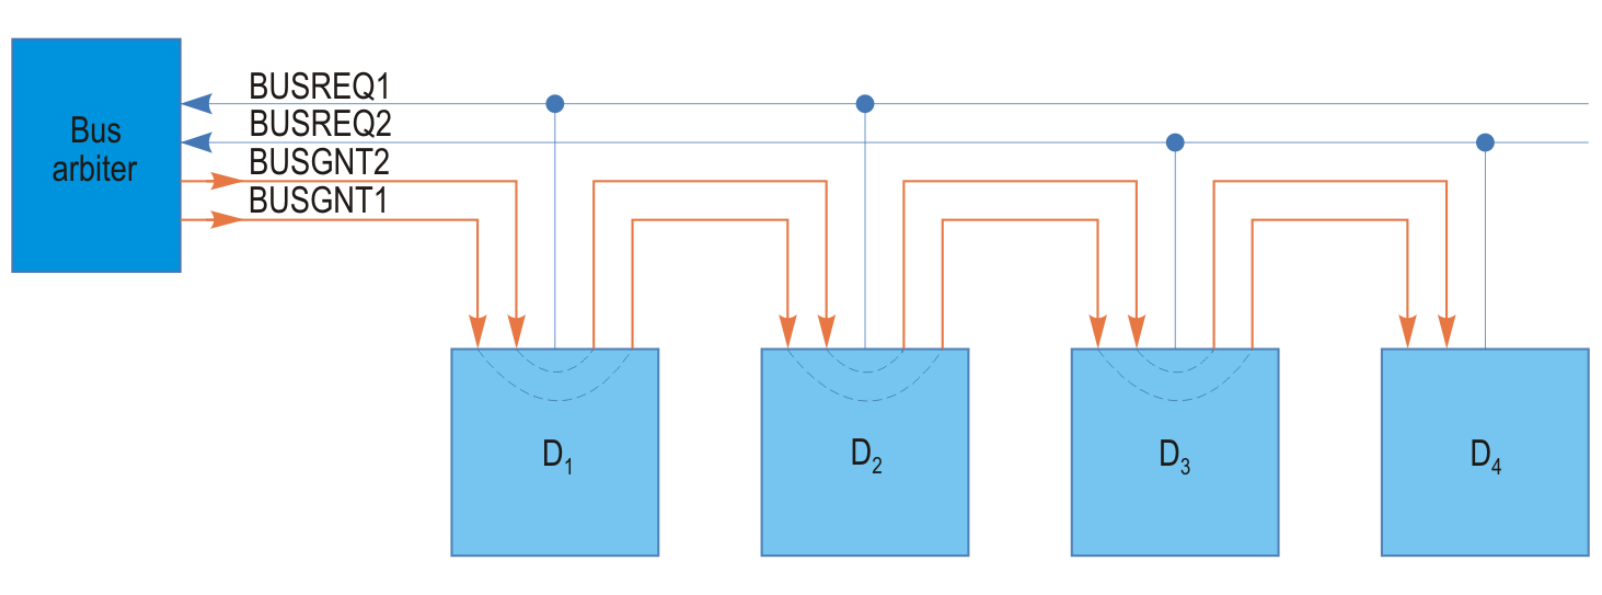

- Each device attaches to one of the bus request lines, according to the device priority

Centralized Bus Arbitration (5)

- Daisy-chaining – Advantages:
  - Small number of control lines required
  - Possibility to connect an unlimited number of devices (theoretically)
- Daisy-chaining – Disadvantages:
  - Fixed priorities of devices
  - A high-priority device may lock out a low-priority device
  - Susceptibility to failures of the *BUSGNT* line

Centralized Bus Arbitration (6)

- Centralized arbitration by independent requesting
  - There are separate *BUSREQ* and *BUSGNT* lines for every device
  - The arbiter may immediately identify all devices requesting the bus and may determine their priority
  - Priority of requests is programmable
  - Disadvantage: to control *n* devices, 2*n* *BUSREQ* and *BUSGNT* lines must be connected to the bus arbiter

Centralized Bus Arbitration (7)

- Centralized arbitration by polling
  - The *BUSGNT* line is replaced with a set of poll-count lines
  - Common *BUSREQ* line
  - The bus arbiter generates a sequence of addresses on the poll-count lines
  - Each device compares these addresses to a unique address assigned to that device
    - On a match, the device asserts the *BUSY* signal and connects to the bus

Centralized Bus Arbitration (8)

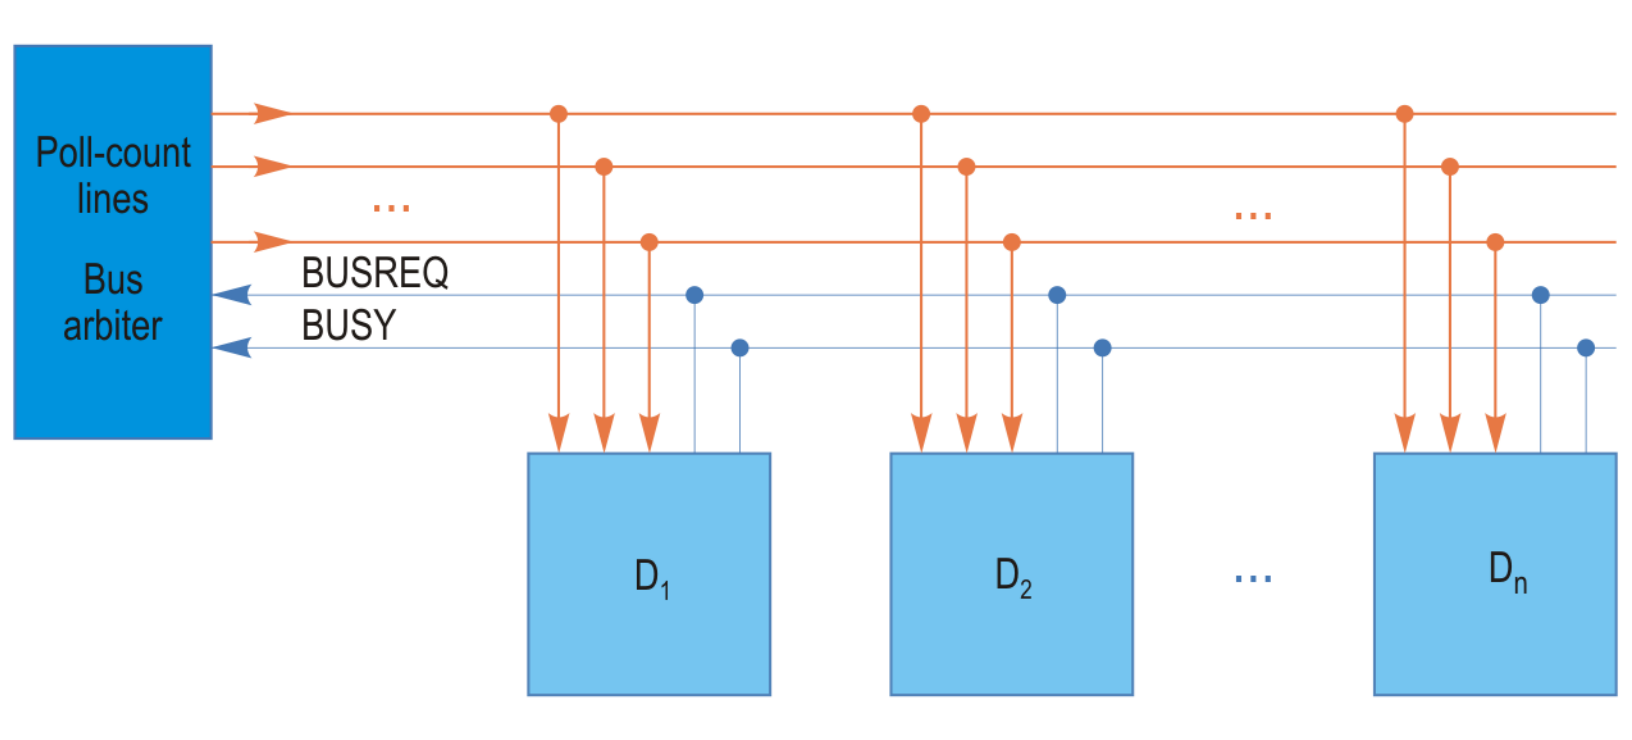

Centralized Bus Arbitration (9)

- The priority of a device is determined by the position of its address in the polling sequence
- Advantages:
  - The sequence can be programmed if the poll-count lines are connected to a programmable register
  - A failure in one device does not affect other devices
- Disadvantage: more control lines are needed

Bus Arbitration

- Bus Arbitration
  - Centralized Bus Arbitration
  - Decentralized Bus Arbitration

Decentralized Bus Arbitration (1)

- There is no bus arbiter
- Example of decentralized arbitration
  - *n* prioritized bus request lines $\rightarrow$ *n* devices
  - To use the bus, a device asserts its request line
  - All devices monitor all the request lines
  - Disadvantages: more bus lines required; the number of devices is limited

Decentralized Bus Arbitration (2)

- Example of decentralized arbitration with only three lines
  - *BUSREQ* $\rightarrow$ wired OR
  - *BUSY* $\rightarrow$ asserted by the bus master
  - Bus arbitration $\rightarrow$ daisy chained
  - The method is similar to the daisy-chain arbitration, but without using an arbiter
  - Advantages: lower cost; higher speed; not subject to arbiter failure

Decentralized Bus Arbitration (3)

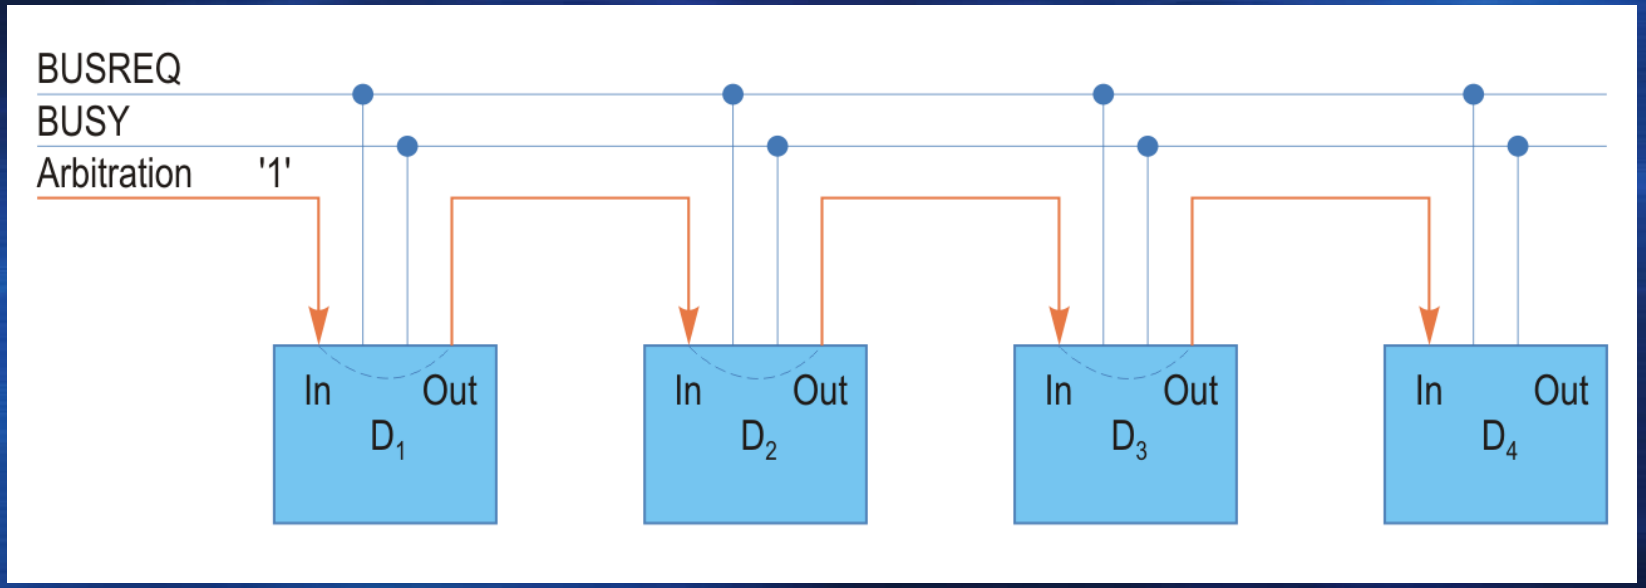

3. Computer Buses

- Introduction
- Electrical Considerations
- Data Transfer Synchronization
- Parallel and Serial Buses
- Bus Arbitration
- PCI Bus
- PCI Express Bus
- Other Serial Buses
- VME Bus

PCI Bus

- PCI Bus
  - PCI Bus Overview
  - PCI Bus Operation
  - PCI-X Bus

PCI Bus Overview (1)

- PCI - *Peripheral Component Interconnect*
- Originally developed by Intel
  - Initial intention: standard for interconnecting the fast circuits on the motherboard
- First version (1.0)
  - Defined mandatory design rules
  - Signals and connections were not defined
- Later on, detailed electrical and functional specifications have been defined for the bus

PCI Bus Overview (2)

- Version 2.0: 33 MHz, up to 132 MB/s
- Other versions: 2.1, 2.2, 2.3, 3.0
- Optional extensions:
  - 64 bits or 66 MHz: up to 264 MB/s
  - 64 bits and 66 MHz: up to 528 MB/s
- The PCI specifications are updated by the *PCI Special Interest Group* (PCI-SIG)
- The PCI bus is not specific to Intel processors

PCI Bus Overview (3)

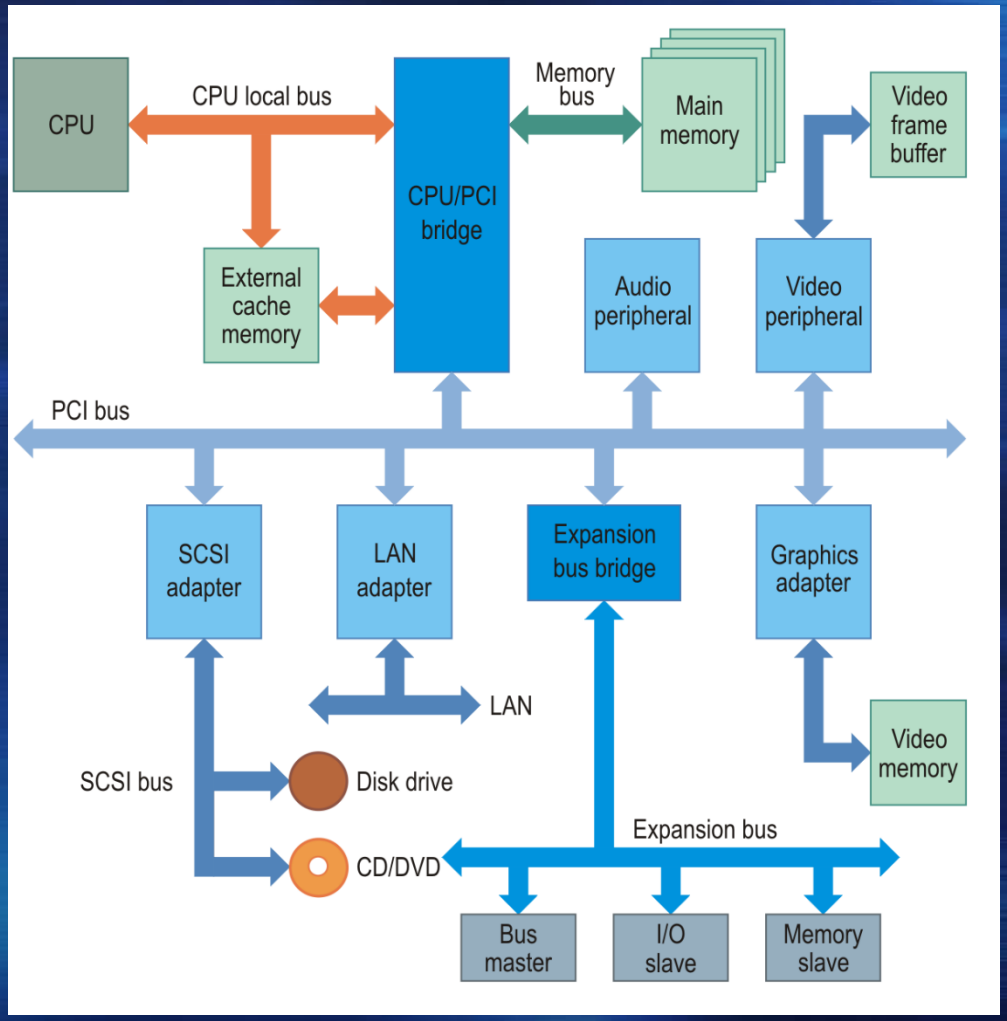

PCI Bus Overview (4)

- The specifications impose a limit of 10 electrical loads (3 expansion boards)
  - Can be extended with PCI-to-PCI bridges
- PCI expansion boards are configured automatically for bus transactions
- PCI devices implement a set of configuration registers (64 x 32 bits)
  - Contain information about presence of the device, device type, address space required
  - The software configures the device's memory and I/O address decoders

PCI Bus

- PCI Bus
  - PCI Bus Overview
  - PCI Bus Operation
  - PCI-X Bus

PCI Bus Operation (1)

- Synchronous operation
- Multiplexed address and data lines, AD
- A centralized bus arbitration is used
  - Performed while another initiator controls the bus → hidden arbitration
- PCI Bus Transactions
  - Performed between a *master* (initiator) and a *slave* (target) device
  - Consist of an address phase followed by one or more data phases

PCI Bus Operation (2)

- Most PCI transactions are performed in burst mode
- A burst transfer consists of:
  - A single address phase
  - Multiple data phases
- Another advantage: bus arbitration must be performed only once
- The target device latches the start address and increments it in each data phase

PCI Bus Operation (3)

- The PCI bus does not need terminators
  - Signal reflections do occur
  - Signal reflections are used as an advantage
  - To assert a signal, a device drives the signal line only to half its required level
  - The signal is reflected back and doubled up to the required activation voltage
  - Advantages: reduced current; reduced driver size

PCI Bus Operation (4)

- PCI Bus Interrupts
  - The PCI bus provides four level-sensitive interrupt request lines, INTA# .. INTD#
  - PCI interrupt request lines are shareable
    - The lines use open-drain technology
    - Multiple devices connected to the same line can assert it simultaneously
  - A particular pattern on the control lines indicates an interrupt acknowledge cycle

PCI Bus Operation (5)

- Interrupt routing
  - Connecting the device's PCI INTx# line to a system IRQ line
  - Interrupt routing should be programmable by the software
- The PCI configuration registers store information about the interrupts
  - Interrupt pin register → the interrupt request line that is used by the device
  - Interrupt line register → interrupt routing

PCI Bus

- PCI Bus
  - PCI Bus Overview
  - PCI Bus Operation
  - PCI-X Bus

PCI-X Bus (1)

- Higher-performance extension of the conventional PCI bus
- Ensures the transfer rates required for connections such as Gigabit Ethernet, Fiber Channel, and InfiniBand
- Initially used for servers and workstations
- Two versions standardized by PCI-SIG: version 1.0 and version 2.0

PCI-X Bus (2)

- PCI-X version 1.0
  - Clock rates up to 133 MHz, 32 or 64 bits
  - Improvements of the conventional protocol
    - Split transactions: an initiator makes a request for a transfer and releases the bus
    - Byte count: an initiator specifies in advance the number of bytes requested
  - Compatibility with previous versions
    - Hardware: operation at 33 or 66 MHz
    - Software: at OS, BIOS, and device driver levels

PCI-X Bus (3)

- PCI-X version 2.0
  - Higher clock rates
    - PCI-X 266 (DDR – *Double Data Rate*): 266 MHz
    - PCI-X 533 (QDR – *Quad Data Rate*): 533 MHz
    - PCI-X 1066: 1066 MHz
    - Maximum performance is 64 times higher compared to the first PCI generation
  - The PCI-X 133 and later variants allow to use a single connector, one electrical load → point-to-point applications

PCI-X Bus (4)

- New features:
  - ECC (*Error Correcting Code*): allows to correct one-bit errors
  - New configuration registers
  - Improved protocol: increases the bus utilization and bus efficiency
  - Strobe signals (PCI-X 266 and PCI-X 533): drive the clock inputs of data buffers
  - 1.5-V signals: allow operation at higher frequencies

3. Computer Buses

- Introduction
- Electrical Considerations
- Data Transfer Synchronization
- Parallel and Serial Buses
- Bus Arbitration
- PCI Bus
- PCI Express Bus
- Other Serial Buses
- VME Bus

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

PCI Express Bus Overview (1)

- PCI-E, PCIe
- Originates from initial specifications of the 3GIO (*Third Generation I/O*) interface
  - Later on, the specifications have been transferred to PCI-SIG
- Serial bus: reduced board complexity, lower pin count, lower cost
- Software model compatible with conventional PCI architecture

PCI Express Bus Overview (2)

- Retains the advantageous features of the previous PCI buses:
  - Same communication model
  - Same address spaces
  - Same transaction types
- Introduces various improvements:
  - Serial connection: eliminates the disadvantages of parallel buses → difficulty of synchronization
  - Point-to-point connection

PCI Express Bus Overview (3)

- Packet-based protocol
- Scalable performance → variable number of communication lanes
- Quality of Service (QoS) feature → differentiated performance
- Advanced power management
- Advanced error reporting and handling
- Possibility of connecting and disconnecting the peripheral devices during operation

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

PCI Express Link (1)

- Minimal PCIe link: two unidirectional communication channels
  - Packets are transmitted: data, commands
  - Channel: two wires with differential signals
  - Communication *lane*
- PCIe link with multiple communication lanes: xN
- Link width and frequency of operation: set automatically

PCI Express Link (2)

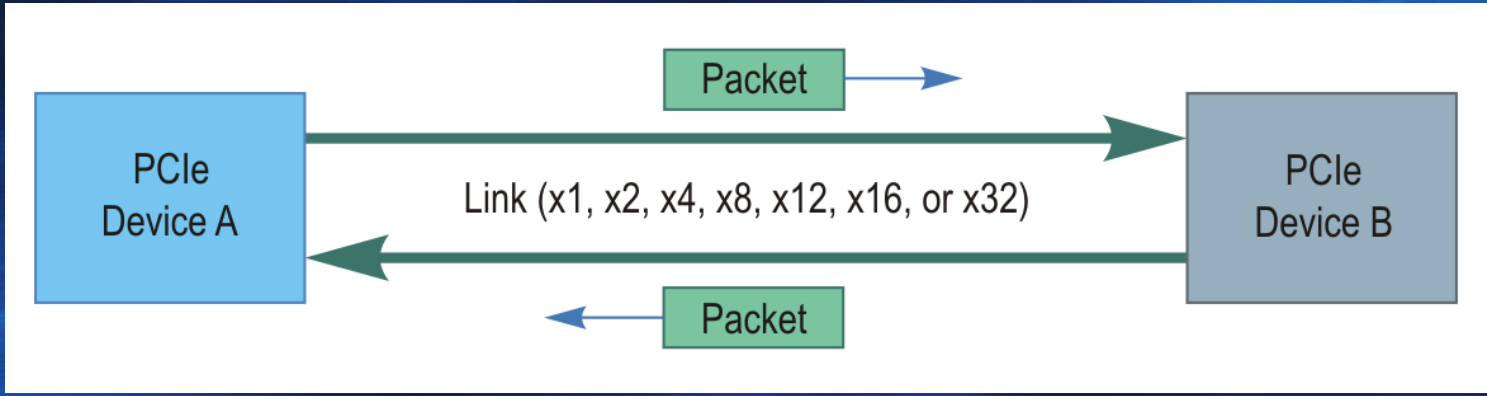

- Operating frequencies:
  - 2.5 GHz (2.5 Gbits/s in each direction, Gen 1)
  - 5 GHz (Gen 2)
  - 8 GHz (Gen 3)
  - 16 GHz (Gen 4)

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

PCI Express Bus Topology (1)

- Root complex → defines a hierarchy
  - Connects the CPU and memory to peripherals
  - PCIe ports: each defines a hierarchy domain
- Endpoints
  - Peripheral devices: initiators (requesters), targets (completers)
  - Up to 8 logical functions (0 .. 7)
- Switch
  - Replaces the shared bus
  - Enables direct communication between two devices

PCI Express Bus Topology (2)

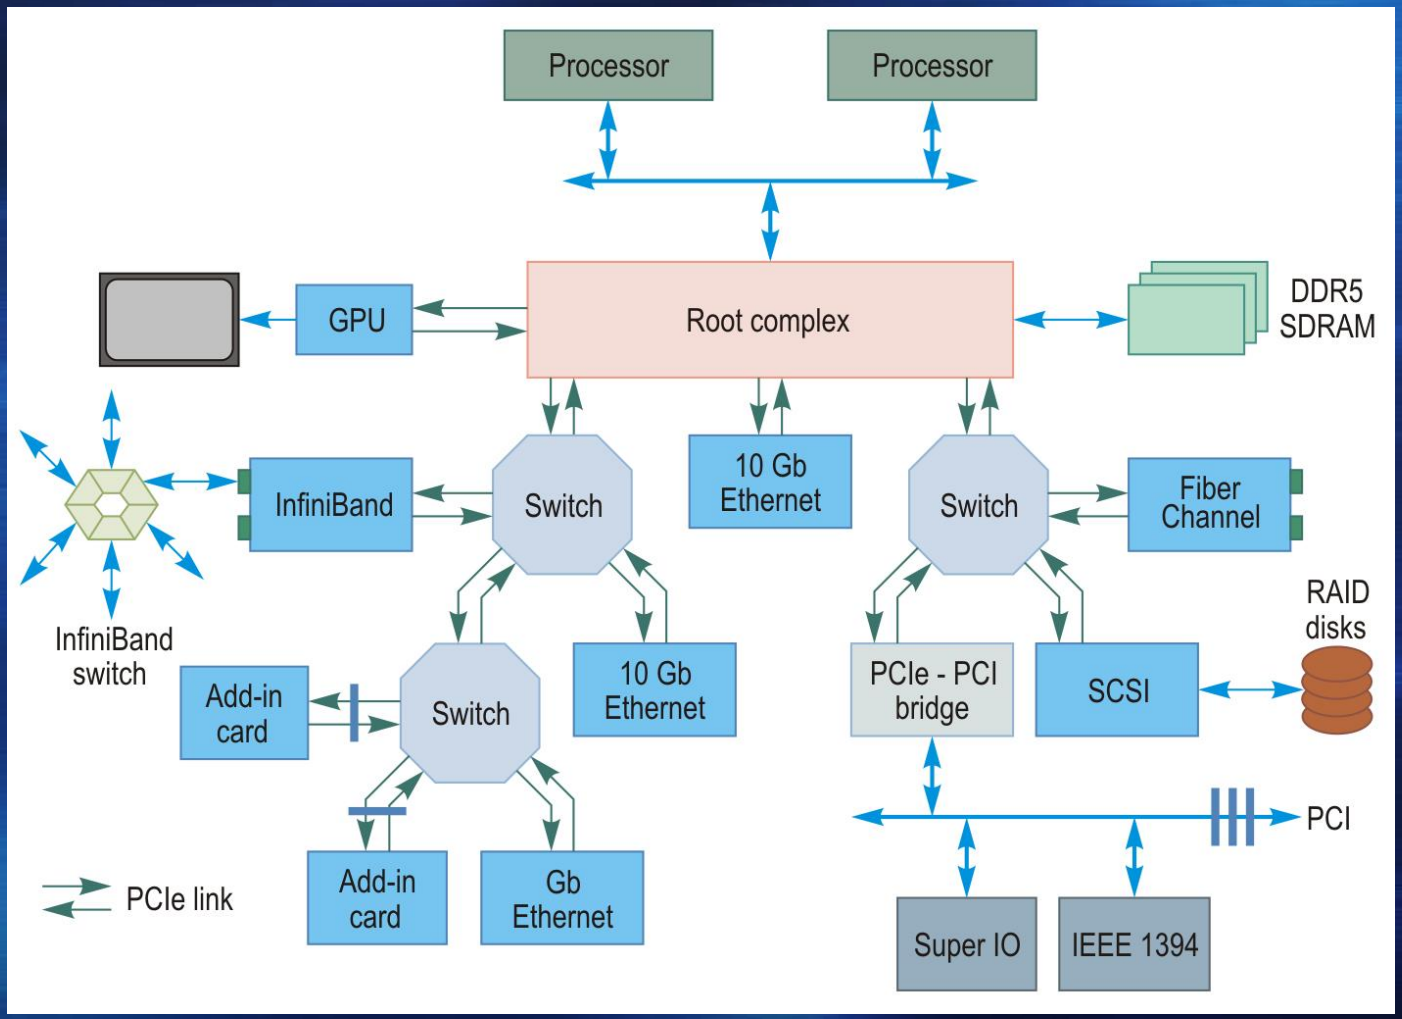

PCI Express Bus Topology (3)

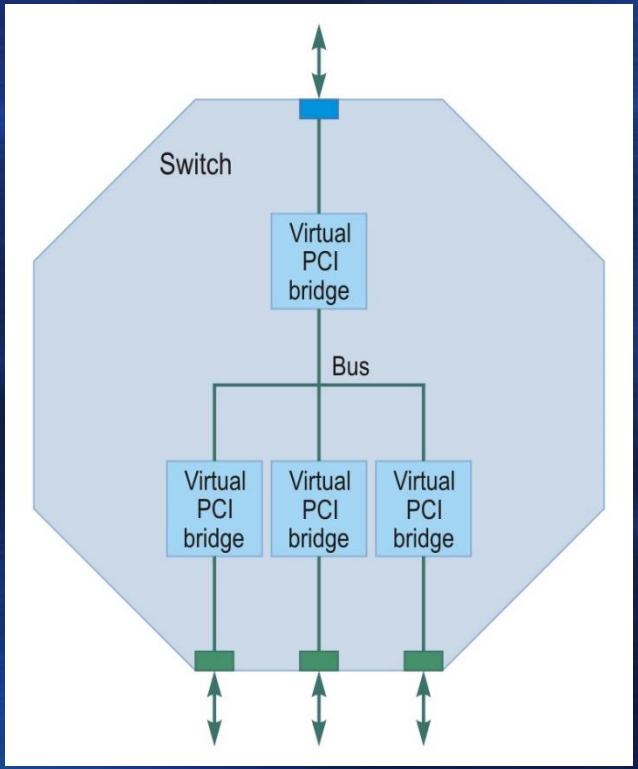

- Switch: assembly of virtual bridges between distinct PCI buses

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

Architecture Layers (1)

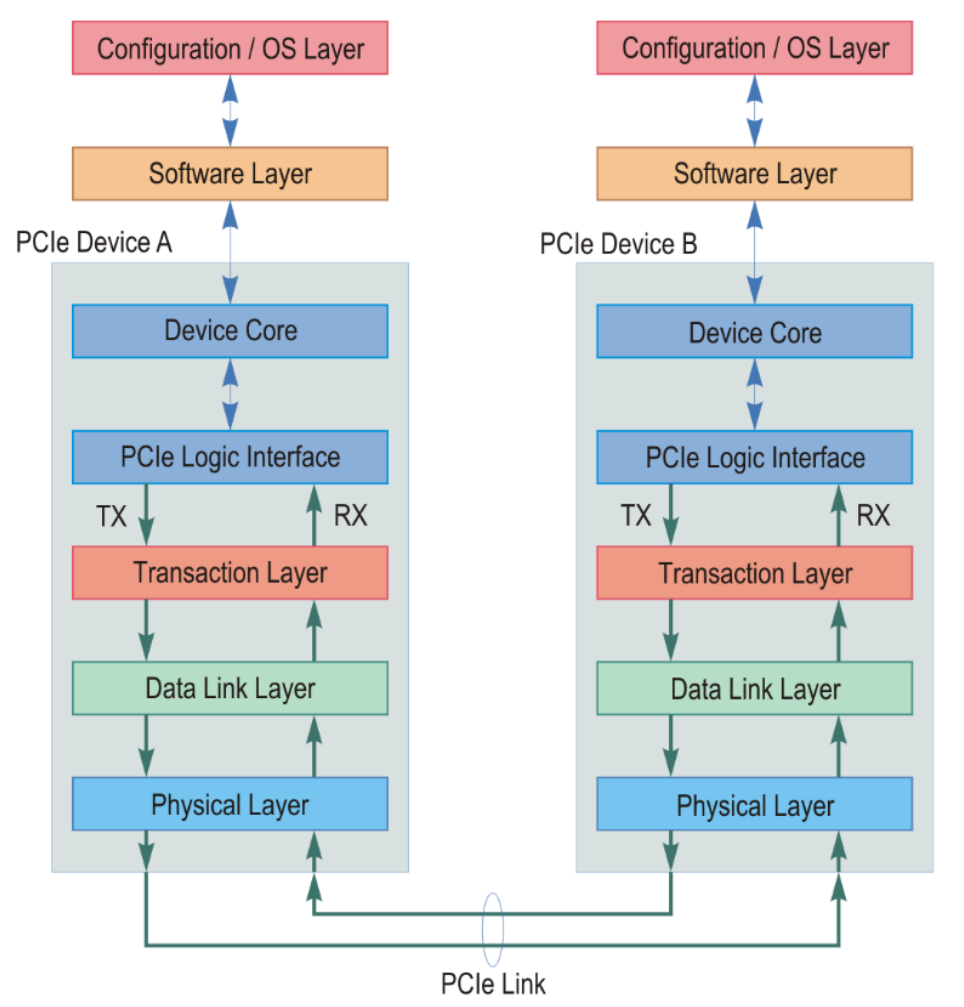

Architecture Layers (2)

- Architecture of PCIe devices
  - The last three layers of PCIe architecture
  - Two sections in each layer: for transmitting and receiving information
- Example: transmit section
  - Transaction layer: forms a packet
  - Data link layer: extends the packet with information for error detection
  - Physical layer: encodes the packet and transmits it via differential signals

Architecture Layers (3)

- Physical layer
  - Data encoding enables to generate a receive clock signal
  - Up to version 3.0: 8b/10b encoding → the bandwidth is reduced with 20%
  - Versions 3.0, 4.0, 5.0: 128b/130b encoding
  - When the link contains several lanes, the bytes are sent interleaved across the lanes
  - Successive bytes are sent on successive lanes → the receive latency is reduced

Architecture Layers (4)

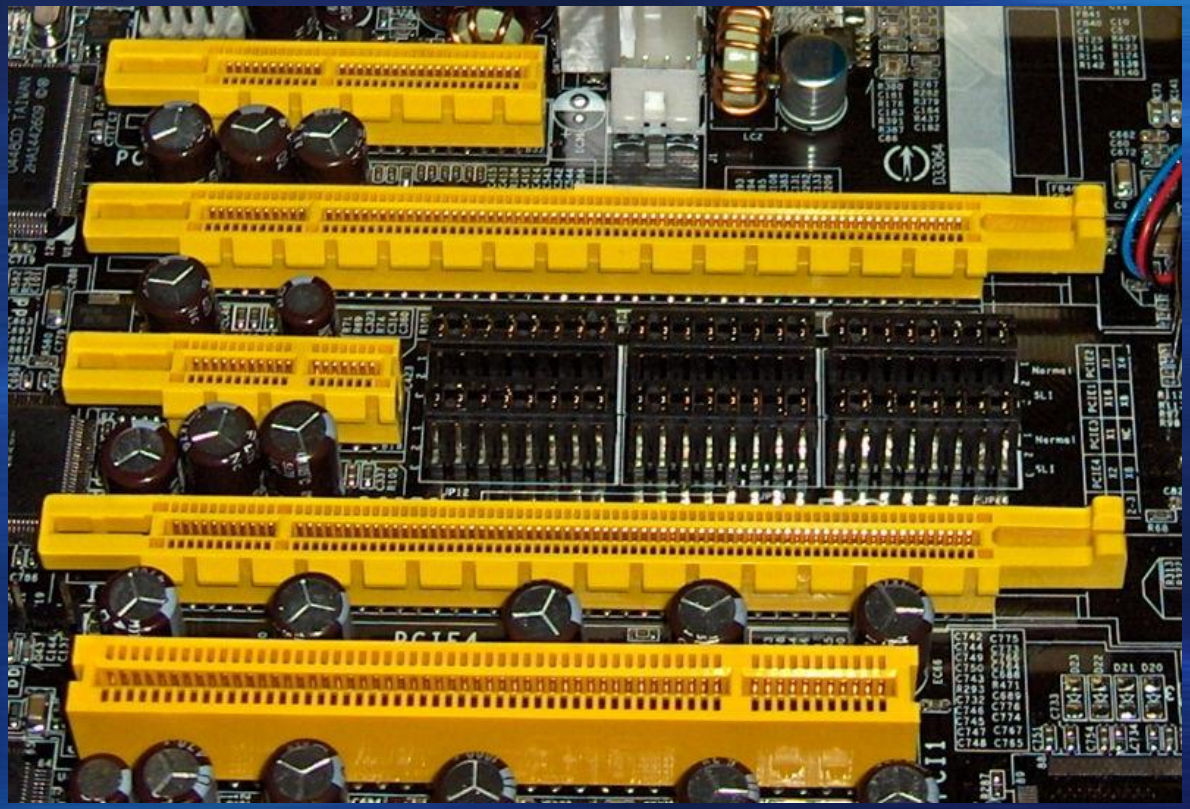

- PCI Express connectors (x4, x16, x1, x16)
- PCI connector (32 bits)

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

PCI Express Transactions (1)

- Transaction: one or more packet transmissions required for a transfer
- Categories of transactions:
  - Memory
  - I/O
  - Configuration
  - Message: power management, interrupt and error signaling

PCI Express Transactions (2)

- Non-posted transactions: the target device returns a completion packet
  - Executed according to the protocol defined for split transactions (PCI-X)
  - The target device stores the information and signals a delayed response
- Posted transactions: the target device does not return a completion packet
  - The time to complete the transaction is reduced

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

PCI Express Interrupts (1)

- Interrupt requests can be signaled in two modes: native mode and legacy mode
- Native mode
  - Message Signaled Interrupts (MSI)
  - Defined as an optional mode for the PCI bus
  - Do not represent PCIe messages, but rather memory write transactions
  - Memory addresses are reserved by the system

PCI Express Interrupts (2)

- Legacy mode
  - Legacy devices use the INTx# interrupt request signals
  - The PCIe bus does not provide the INTx# interrupt lines
  - Special messages are used that act as virtual INTx# wires (e.g., INTA# assertion message)
  - The messages are targeted to the interrupt controller located within the root complex

PCI Express Bus

- PCI Express Bus
  - PCI Express Bus Overview
  - PCI Express Link
  - PCI Express Bus Topology
  - Architecture Layers
  - PCI Express Transactions
  - PCI Express Interrupts
  - Versions of PCI Express Standards

Versions of PCI Express Standards

- Versions 1.0, 1.1 (2002, 2005): 2.5 GT/s
- Versions 2.0, 2.1 (2007): 5 GT/s
- Version 3.0 (2010)
  - 8 GT/s; x16 connector: 15.7 GB/s
- Version 4.0 (2017)
  - 16 GT/s; x16 connector: 31.5 GB/s
- Version 5.0 (2019)
  - 32 GT/s; x16 connector: 63 GB/s
- Version 6.0 (2022)
  - 64 GT/s; x16 connector: 121 GB/s

Summary (1)

- Parallel buses are difficult to improve
  - Main reason: timing skew
- Serial buses replaced most of the parallel buses
  - They have several advantages over parallel buses
- Bus arbitration methods can be centralized or decentralized
  - Centralized arbitration methods: daisy-chaining; independent requesting; polling

Summary (2)

- The PCI bus has been successful for personal computers
  - Its transfer rate is limited due to the parallel nature of the bus
  - Configuration registers enable automatic configuration of PCI devices
  - On the PCI bus, signal reflections are used as an advantage
- The PCI-X bus improves the performance of the parallel PCI bus

Summary (3)

- The PCI Express (PCIe) bus maintains software compatibility with the PCI architecture
  - Introduces high-speed serial connections
  - Other improvements: point-to-point connection; packet-based protocol; scalable performance; QoS feature
  - Topological elements of the PCIe bus: root complex, endpoints, switch
  - Interrupts can be signaled in native mode and legacy mode

Concepts, Knowledge (1)

- Parallel and serial buses
- Timing skew effect
- Advantages of serial buses
- Centralized bus arbitration by daisy-chaining
- Centralized bus arbitration by independent requesting
- Centralized bus arbitration by polling
- Decentralized bus arbitration

Concepts, Knowledge (2)

- PCI bus overview
- PCI bus terminators
- PCI bus arbitration
- PCI bus transactions
- PCI bus interrupts
- Improvements introduced by PCI-X bus version 1.0
- Improvements introduced by PCI-X bus version 2.0

Concepts, Knowledge (3)

- Improvements introduced by the PCIe bus
- PCIe link
- Elements of PCIe bus topology
- Physical layer of the PCIe architecture
- Categories and types of PCIe transactions
- Native mode of PCIe interrupt requests
- Legacy mode of PCIe interrupt requests