

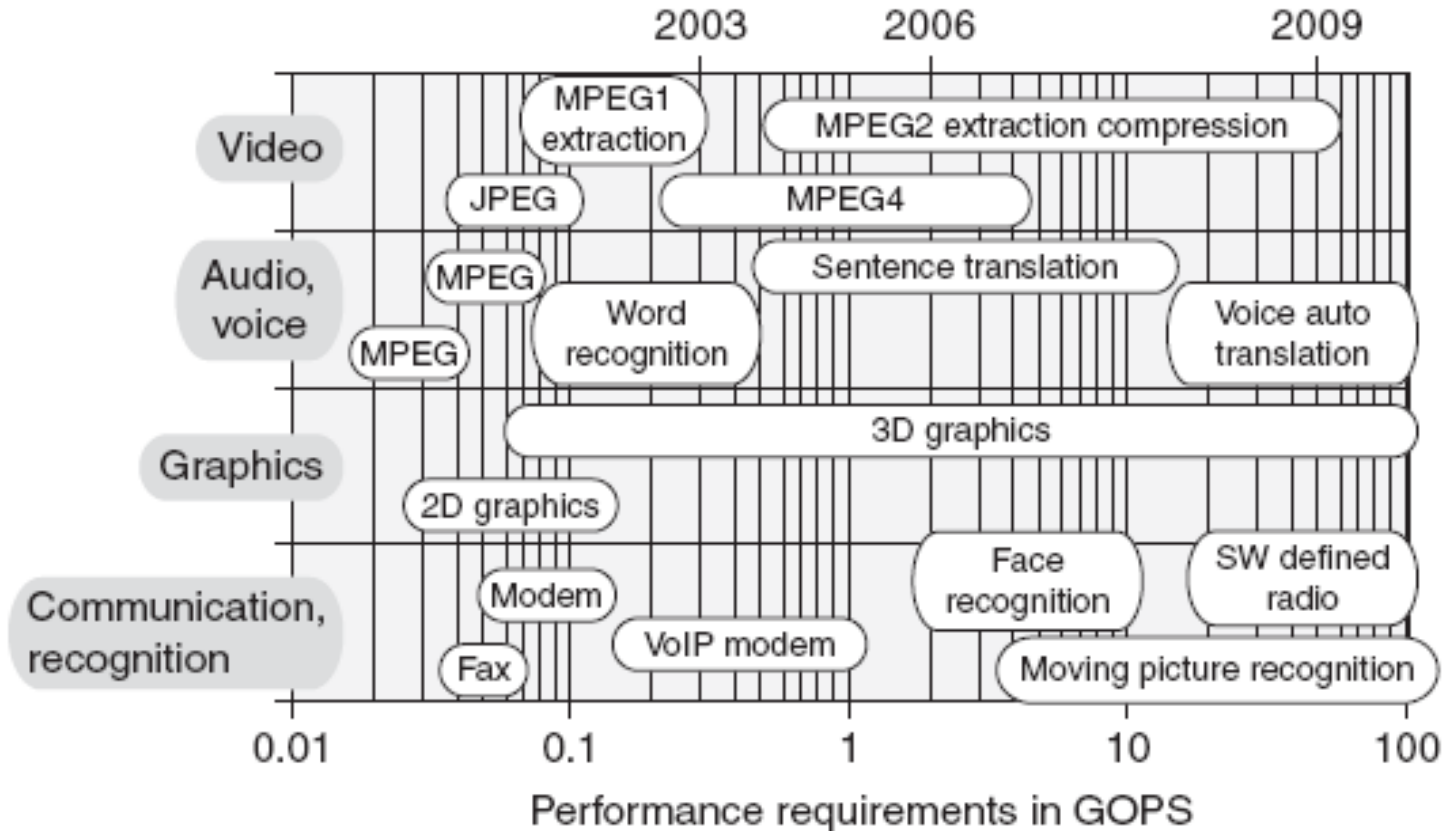
ECE 551

System on Chip Design

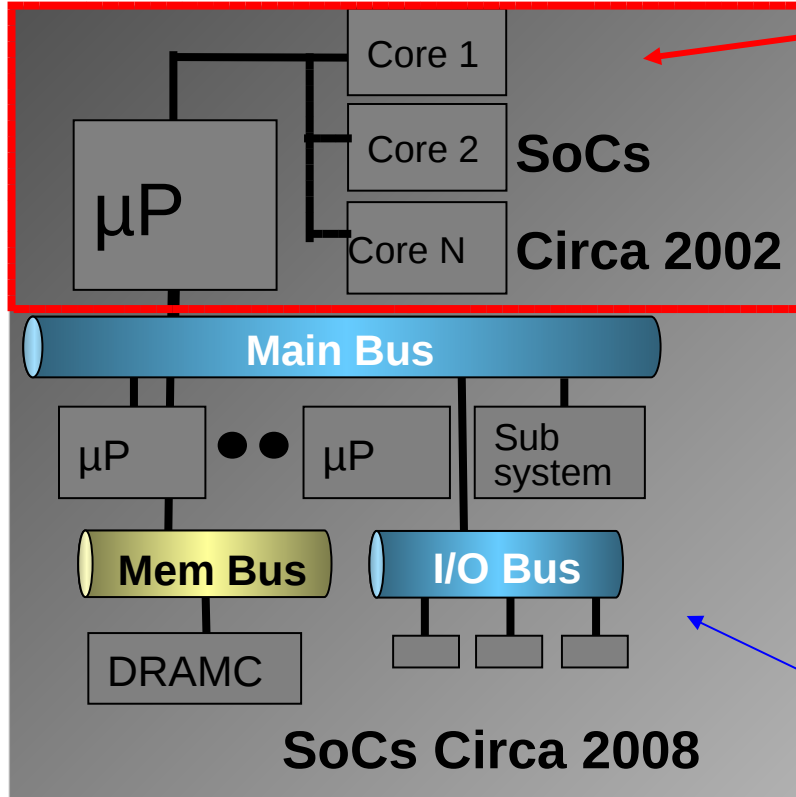
Introducing Bus Communications

Garrett S. Rose
Fall 2018

Emerging Applications Requirements



Data Flow vs. Processing



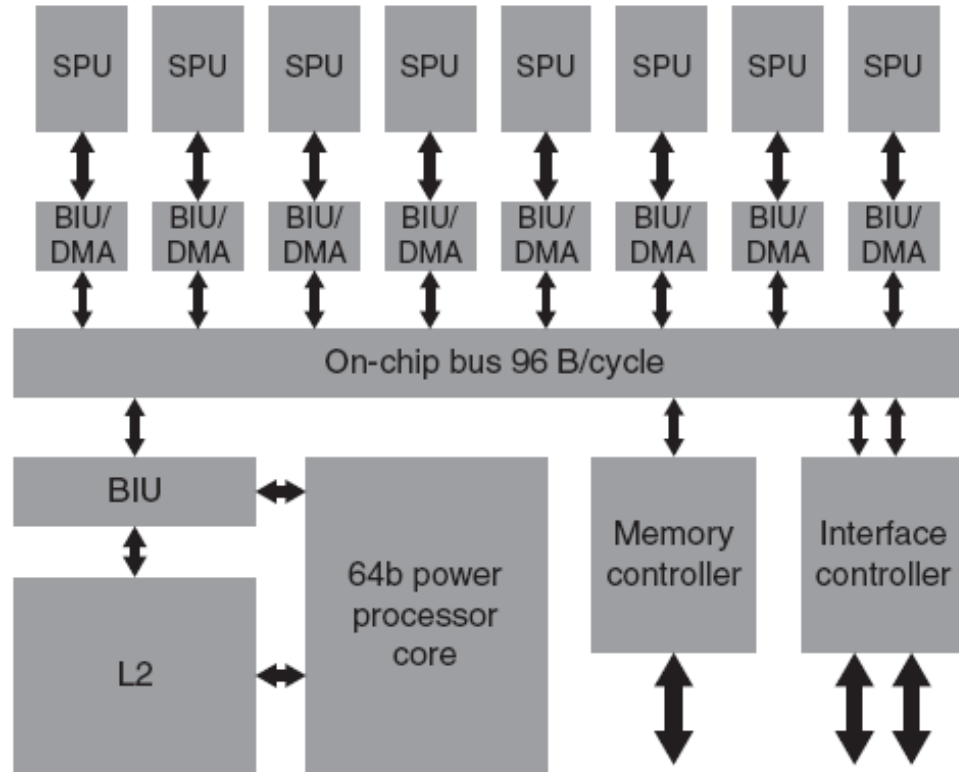
Critical Decision Was uP Choice

- Data flow replacing data processing major design challenge
- Exploding core counts requiring more advanced interconnects
- EDA cannot solve this architectural problem easily
- Complexity too high to hand craft (and verify!)


Critical Decision Is Interconnect Choice

Communication Architecture Design and Verification becoming Highest Priority in Contemporary SoC Design!

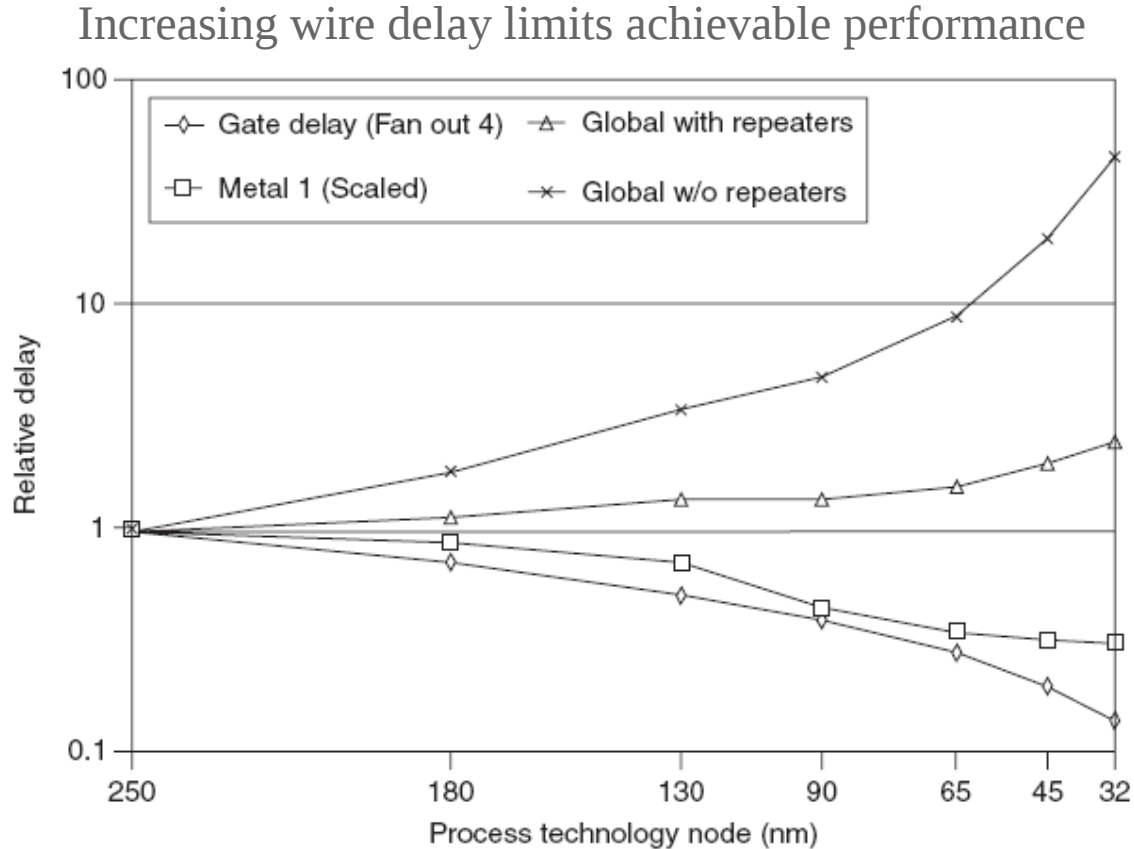
Example: IBM Cell Ring Bus



Tech Scaling Trends: On Chip Interconnect Length

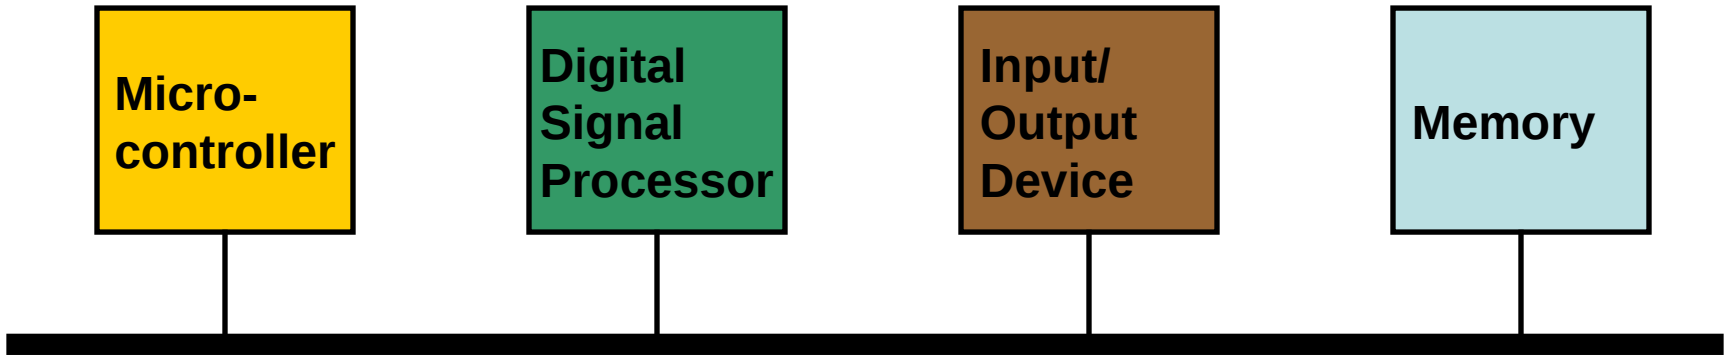


Tech Scaling Trends: Interconnect Performance

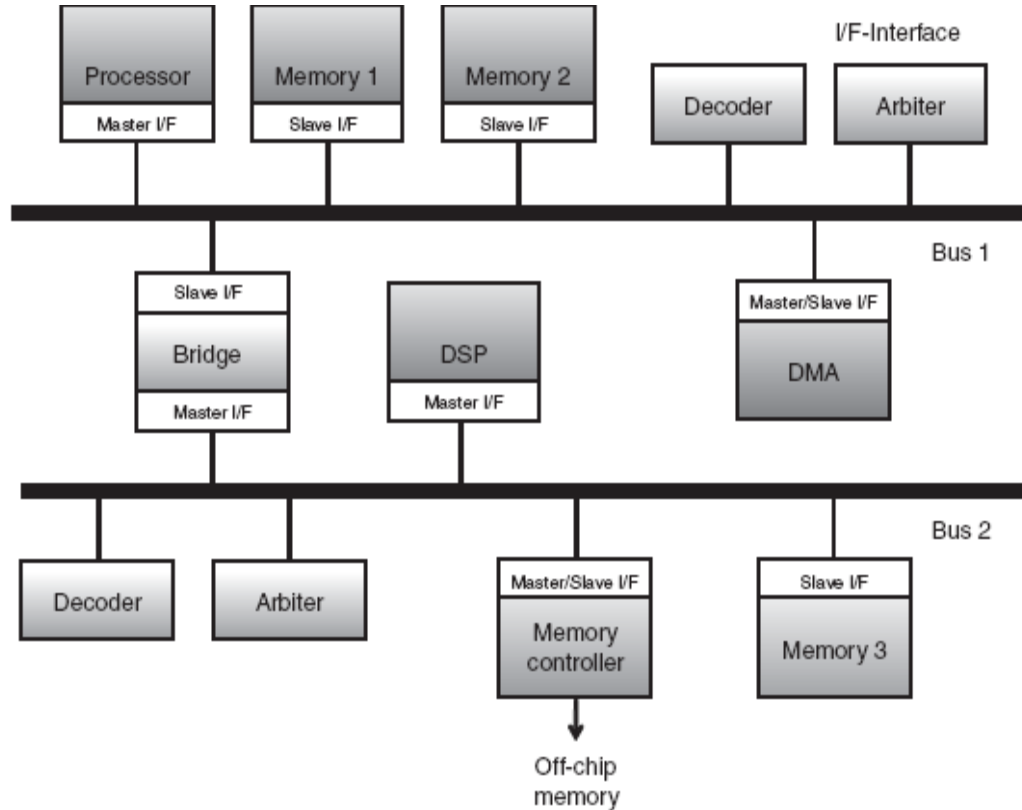


Bus Based Communication Architectures

- Buses are the simplest and most widely used SoC interconnection networks
- Bus – collection of signals (wires) with one or more IP components connected
- Only one IP component can transfer data on the shared bus at any given time




Bus Terminology



Bus Terminology

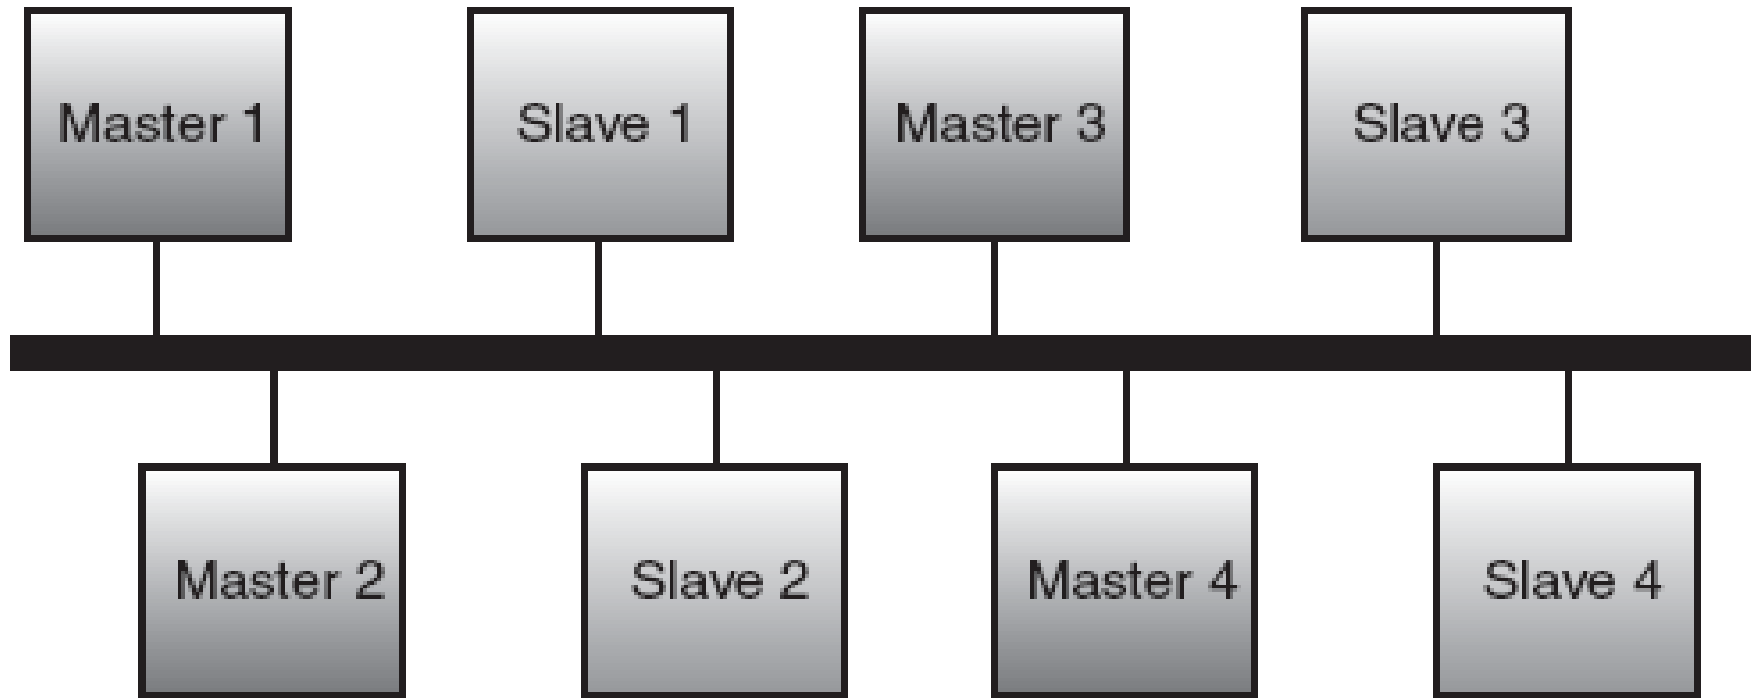
- Master (or Initiator)
 - IP component that initiates a read or write data transfer
- Slave (or Target)
 - IP component that does not initiate transfers and only responds to incoming transfer requests
- Arbiter
 - Controls access to the shared bus
 - Uses arbitration scheme to select master to grant access to bus
- Decoder
 - Determines which component a transfer is intended for
- Bridge
 - Connects two busses
 - Acts as *slave* on one side and *master* on the other

Bus Signal Lines

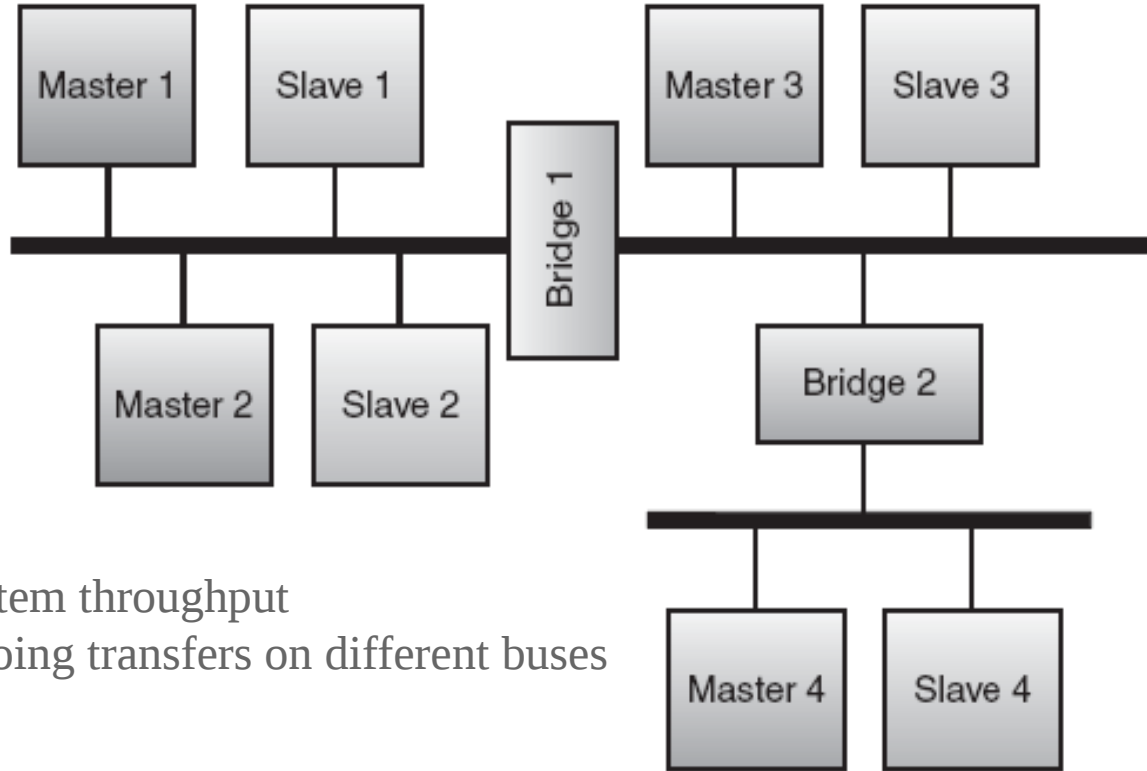


- Address
 - Carry address of destination for which transfer is initiated
 - Can be shared or separate for read, write data
- Data
 - Carry information between source and destination components
 - Can be shared or separate for read, write data
 - Choice of data width critical for application performance
- Control (Requests and acknowledgements)
 - Specify more information about type of data transfer
 - Byte enable, burst size, cacheable/bufferable, write-back/through

Bus Topologies: Shared




Bus Topologies: Hierarchical



- Improves system throughput
- Multiple ongoing transfers on different buses


Bus Topologies: Split Bus



- Reduces impact of capacitance across two segments
- Reduces contention and energy


Bus Topologies:

Full Matrix/Crossbar




Bus Topologies:

Full Matrix/Crossbar



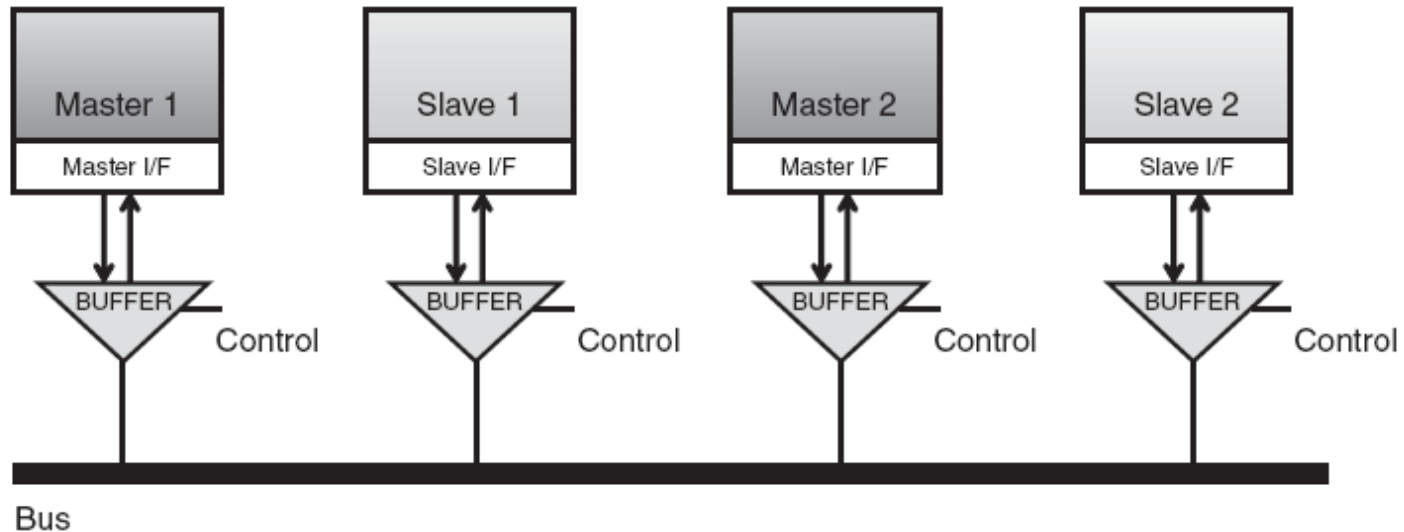
Bus Topologies: Ring Bus

** The IBM CELL Processor uses a ring bus

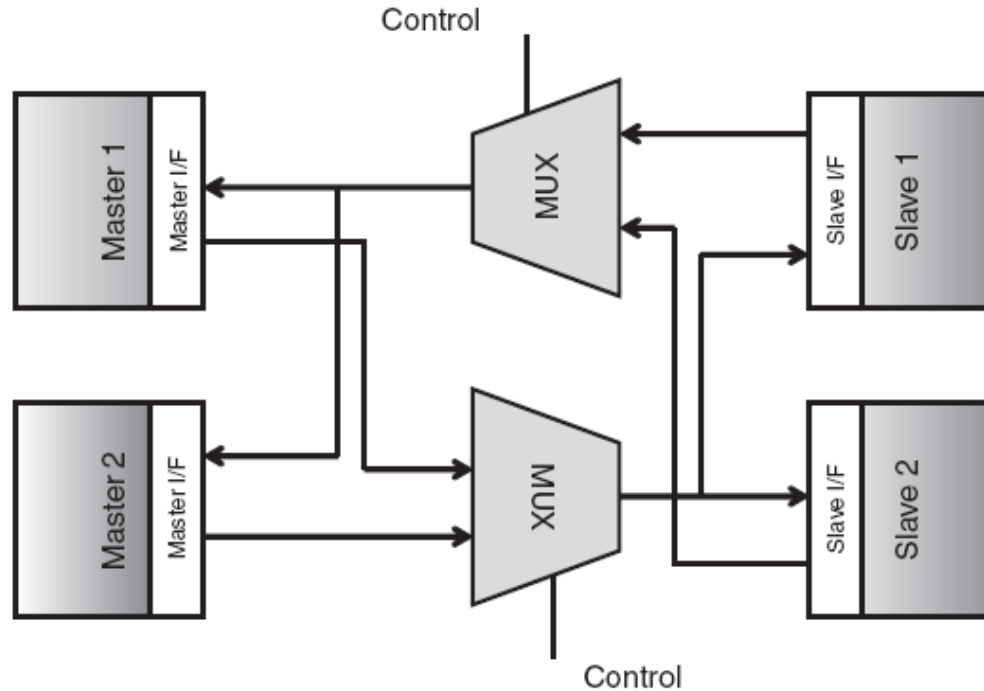


Bus Physical Structure: Tri-State Buffer Based

- Commonly used in off-chip/backplane buses
 - Pro: take up fewer wires, smaller area footprint
 - Con: higher power consumption & delay, hard to debug

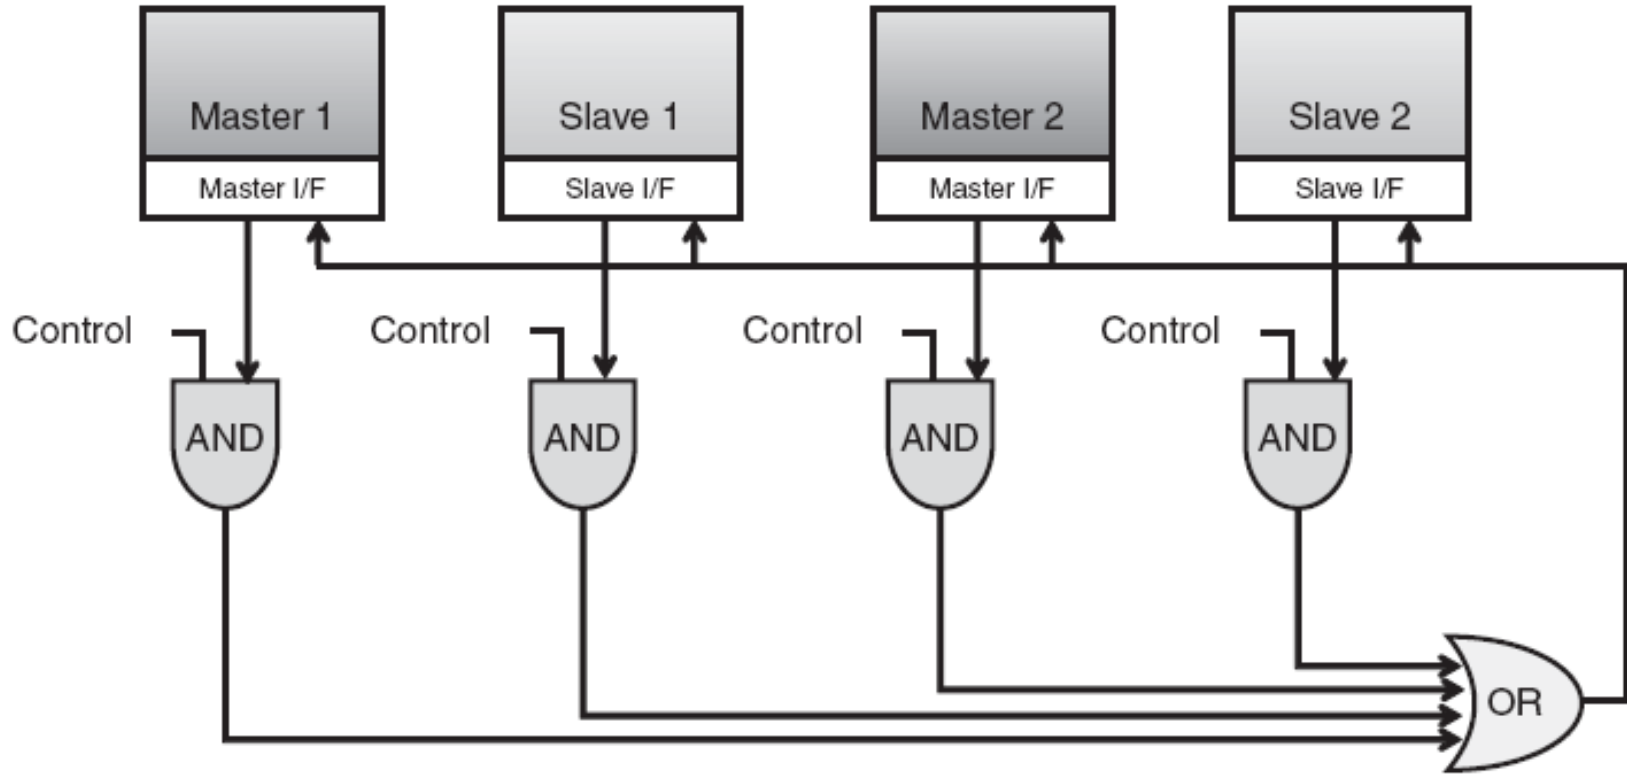


Bus Physical Structure: MUX Based



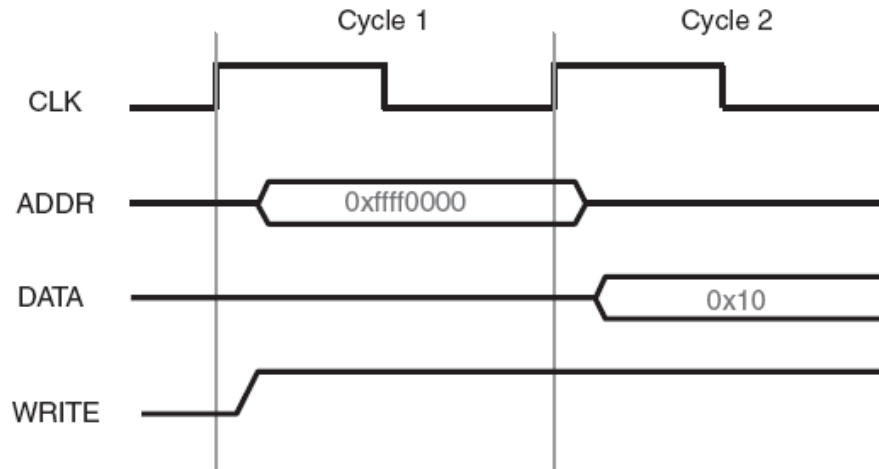
Separate read, write channels

Bus Physical Structure: AND-OR Based




Bus Clocking: Synchronous

- Includes a clock in control lines
- Fixed protocol for communication that is relative to clock
- Involves very little logic and can run very fast
- Require frequency converters across frequency domains



Bus Clocking: Asynchronous

- Not clocked
- Requires a handshaking protocol
 - Performance not as good as that of synchronous bus
 - No need for frequency converters, but does need extra lines
- Does not suffer from clock skew like the synchronous bus

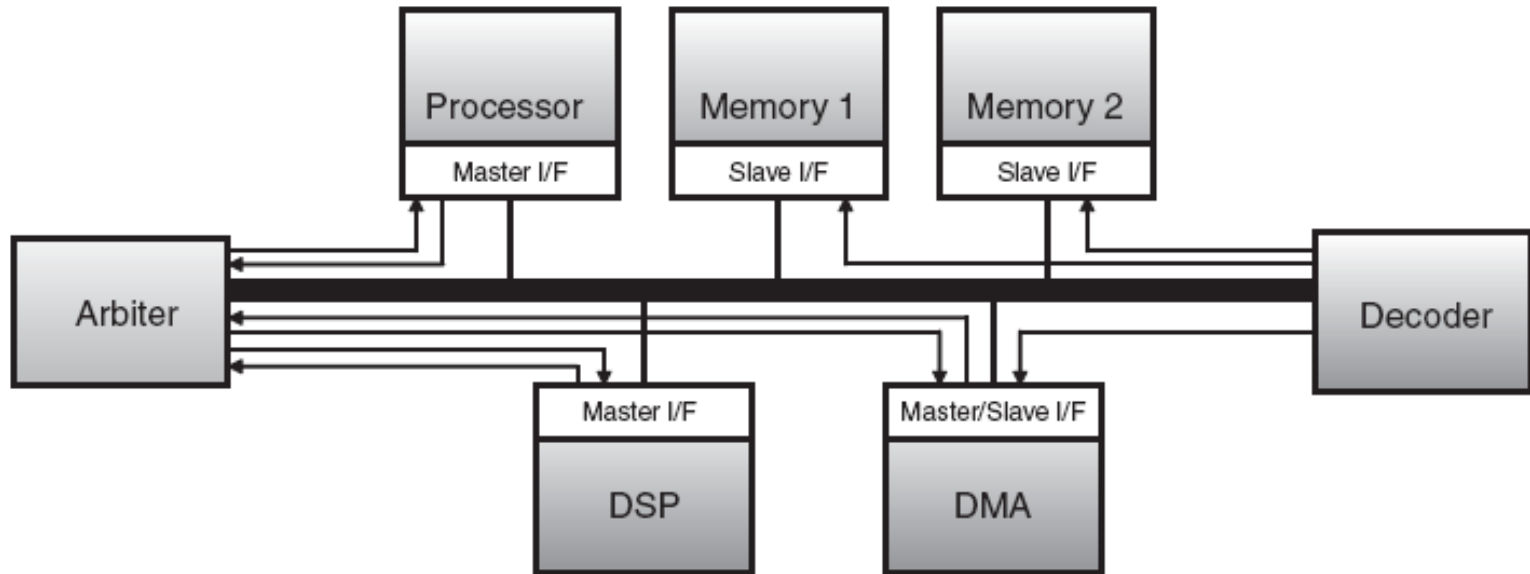


Decoding and Arbitration

- Decoding
 - determines the target for any transfer initiated by a master
- Arbitration
 - decides which master can use the shared bus if more than one master request bus access simultaneously
- Decoding and Arbitration can be:
 - *centralized*
 - *distributed*

Centralized Decoding and Arbitration

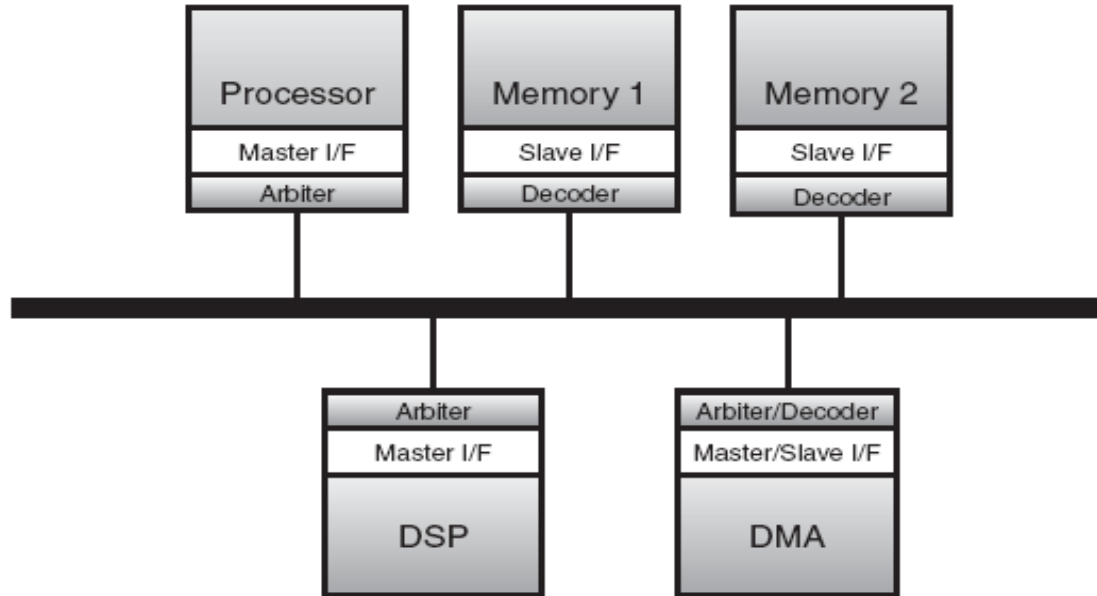
Minimal change is required if new components are added to the system



Distributed Decoding and Arbitration

Pro: requires fewer signals compared to the centralized approach

Con: more hardware duplication, more logic/area, less scalable



Arbitration Schemes

- Random
 - Randomly select master to grant bus access to
- Static priority
 - Masters assigned static priorities
 - Higher priority master request always serviced first
 - Can be pre-emptive (AMBA2) or non-preemptive (AMBA3)
 - May lead to starvation of low priority masters
- Round Robin (RR)
 - Masters allowed to access bus in a round-robin manner
 - No starvation – every master guaranteed bus access
 - Inefficient if masters have vastly different data injection rates
 - High latency for critical data streams

Arbitration Schemes


- TDMA
 - Time division multiple access
 - Assign slots to masters based on BW requirements
 - If a master does not have anything to read/write during its time slots, leads to low performance
 - Choice of time slot length and number critical
- TDMA/RR
 - Two-level scheme
 - If master does not need to utilize its time slot, second level RR scheme grants access to another waiting master
 - Better bus utilization
 - Higher implementation cost for scheme (more logic, area)

Arbitration Schemes

- Dynamic priority
 - Dynamically vary priority of master during application execution
 - Gives masters with higher injection rates a higher priority
 - Requires additional logic to analyze traffic at runtime
 - Adapts to changing data traffic profiles
 - High implementation cost (several registers to track priorities and traffic profiles)
- Programmable priority
 - Simpler variant of dynamic priority scheme
 - Programmable register in arbiter allows software to change priority

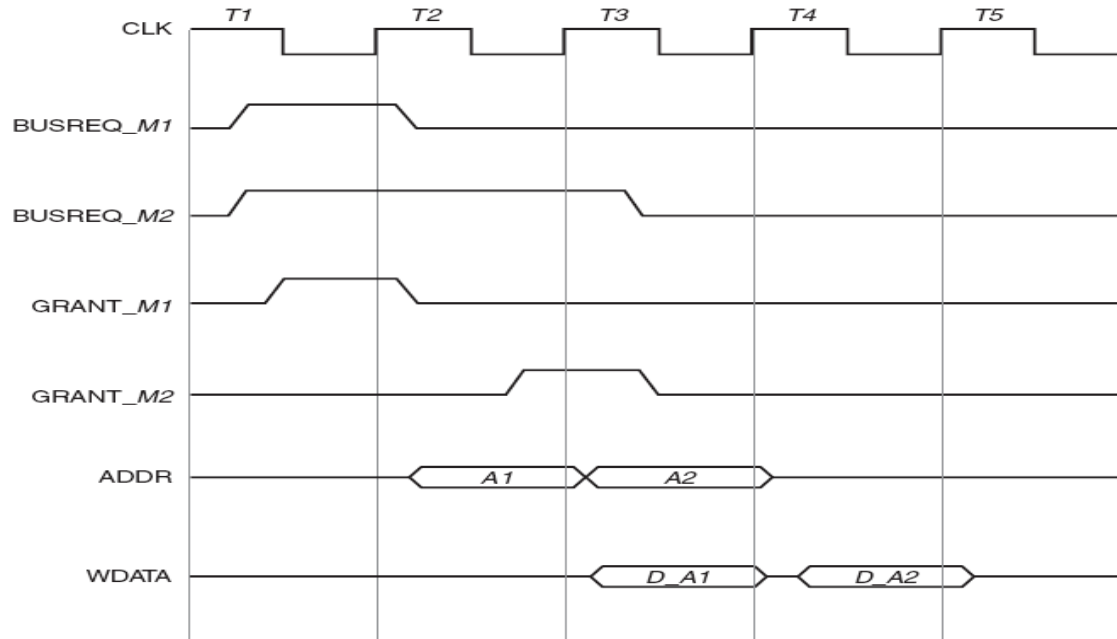
Bus Data Transfer Modes: Single Non-Pipelined

- Simplest transfer mode
- first request for access to bus from arbiter
- on being granted access, set address and control signals
- Send/receive data in subsequent **cycles**



Bus Data Transfer Modes: Pipelined

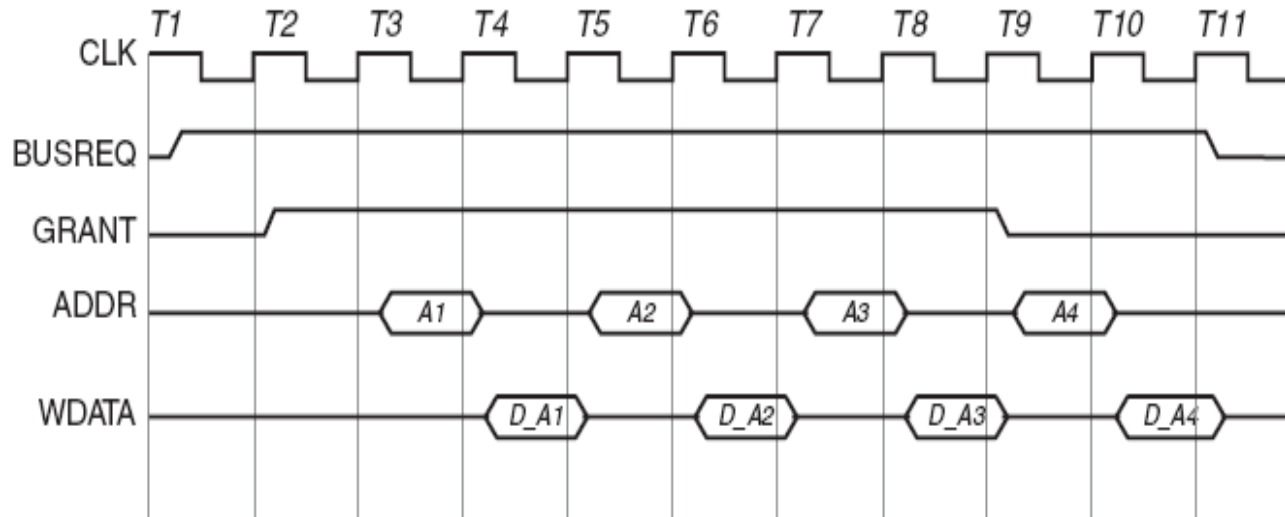
- Overlap address and data phases
- Only works if separate address and data buses are present



Bus Data Transfer Modes:

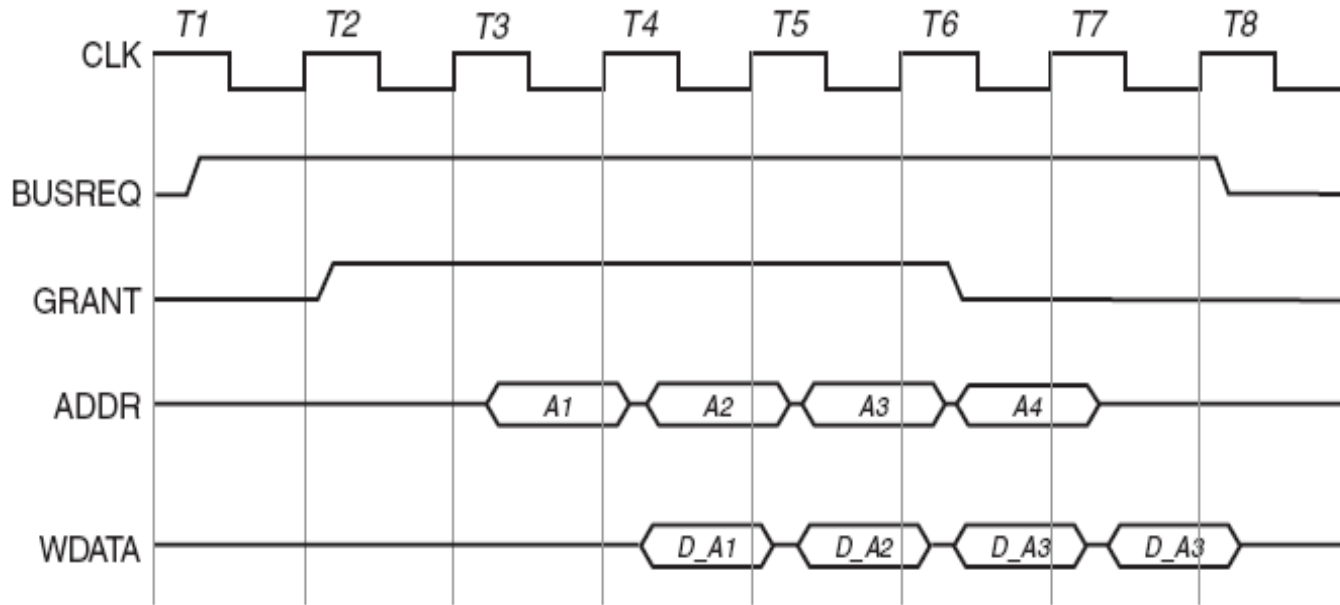
Non-Pipelined Burst

- Send multiple data, arbitrate once for entire transaction
- Master indicates to arbiter intention to perform burst transfer
- Saves time spent requesting for arbitration



Bus Data Transfer Modes: Pipelined Burst

- Useful when separate address and data buses available
- Reduces data transfer latency



Bus Data Transfer Modes: Split Transfer

- If slaves take a long time to read/write data, it can prevent other masters from using the bus
- Split transfers improve performance by ‘splitting’ a transaction
 - Master sends read request to slave
 - Slave relinquishes control of bus as it prepares data
 - Arbiter can grant bus access to another waiting master
 - Allows utilizing otherwise idle cycles on the bus
 - When slave is ready, it requests bus access from arbiter
 - On being granted access, it sends data to master
- Explicit support for split transfers required from slaves and arbiters (additional signals, logic)

Bus Data Transfer Modes:

Out-of-Order Transfer

- Multiple transfers from different masters, or even from the same master, are SPLIT by slave and are in progress simultaneously on single bus
- Masters can initiate transfers without waiting for earlier transfers to complete
- Allows better parallelism, performance in buses
- Additional signals needed to transmit IDs for every data transfer in the system
- Master interfaces need to be extended to handle data transfer IDs and be able to reorder received data
- Slave interfaces have out-of-order buffers for reads, writes, to keep track of pending transactions, plus logic for processing IDs
 - Any application typically has a limited buffer size beyond which performance doesn't increase

Bus Data Transfer Modes: Broadcast

- Every time a data item is transmitted over a bus, it is physically broadcast to every component on the bus
- Useful for snooping and cache coherence protocols
- Example: when several components on bus have a private cache fed from a single memory, a problem arises when the memory is updated
 - when a cache line is written to memory by a component
- It is essential that private caches of the components on the bus invalidate (or update) their cache entries
 - to prevent reading incorrect values
- Broadcasting allows address of the memory location (or cache line) being updated to be transmitted to all the components on the bus, so they can invalidate (or update) their local copies