CSE 520
Computer Architecture II

On-Chip Interconnect Networks

Prof. Michel A. Kinsy

Modern Computer Architecture Components

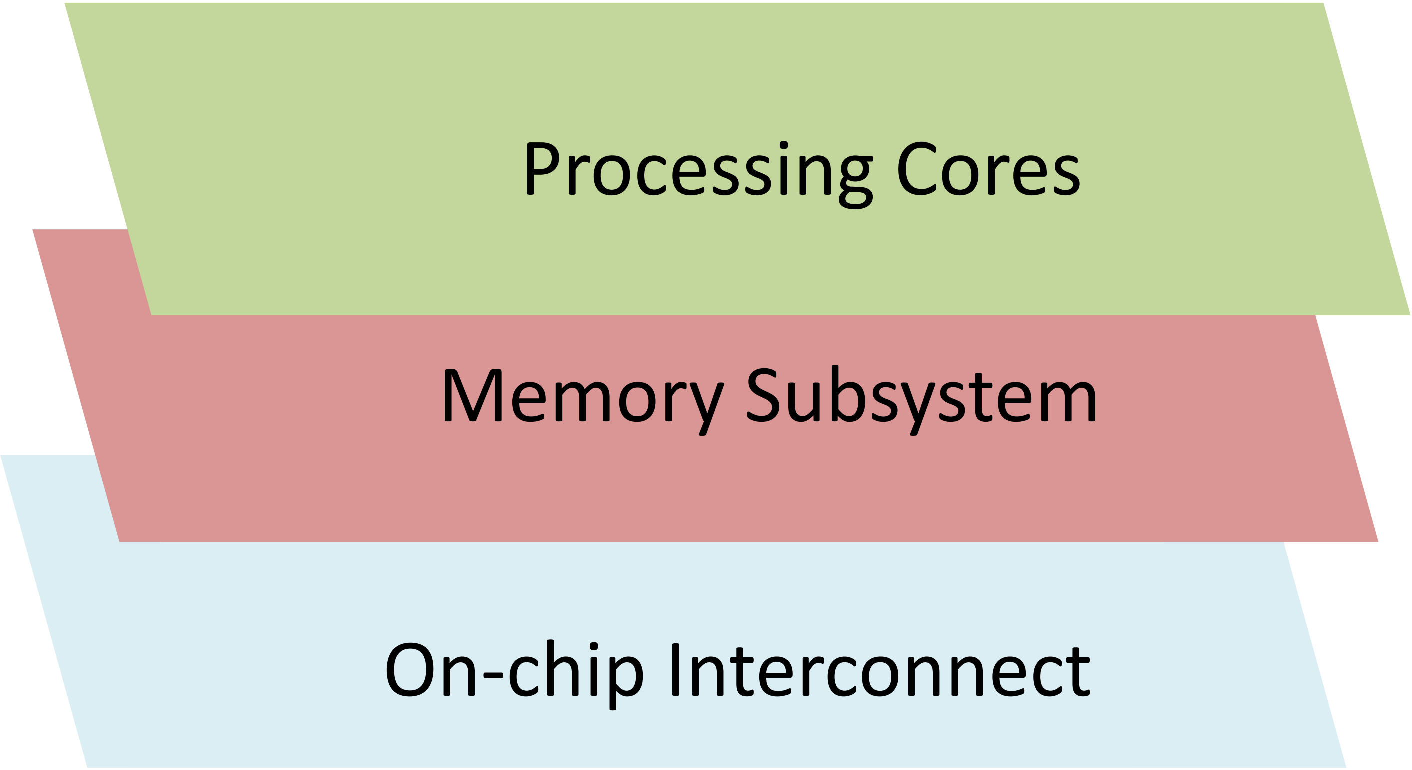

- On-chip network handles cache lines and cache coherence messages between processor cores and memories

Multicore Processors

- Trend: towards ever larger multicore and many-core chip
  - Intel's Single-chip Cloud Computer with 48 cores
  - Tilera Corporation TILE-Gx with 100 cores
- Multi/many-core help overcome diminishing returns of increasingly complex single-core processors
  - Communication is critical to the multicore performance
  - Communications between cores, cache banks, DRAM controllers
  - Delays in information can stall the pipeline

Shared Bus Communication

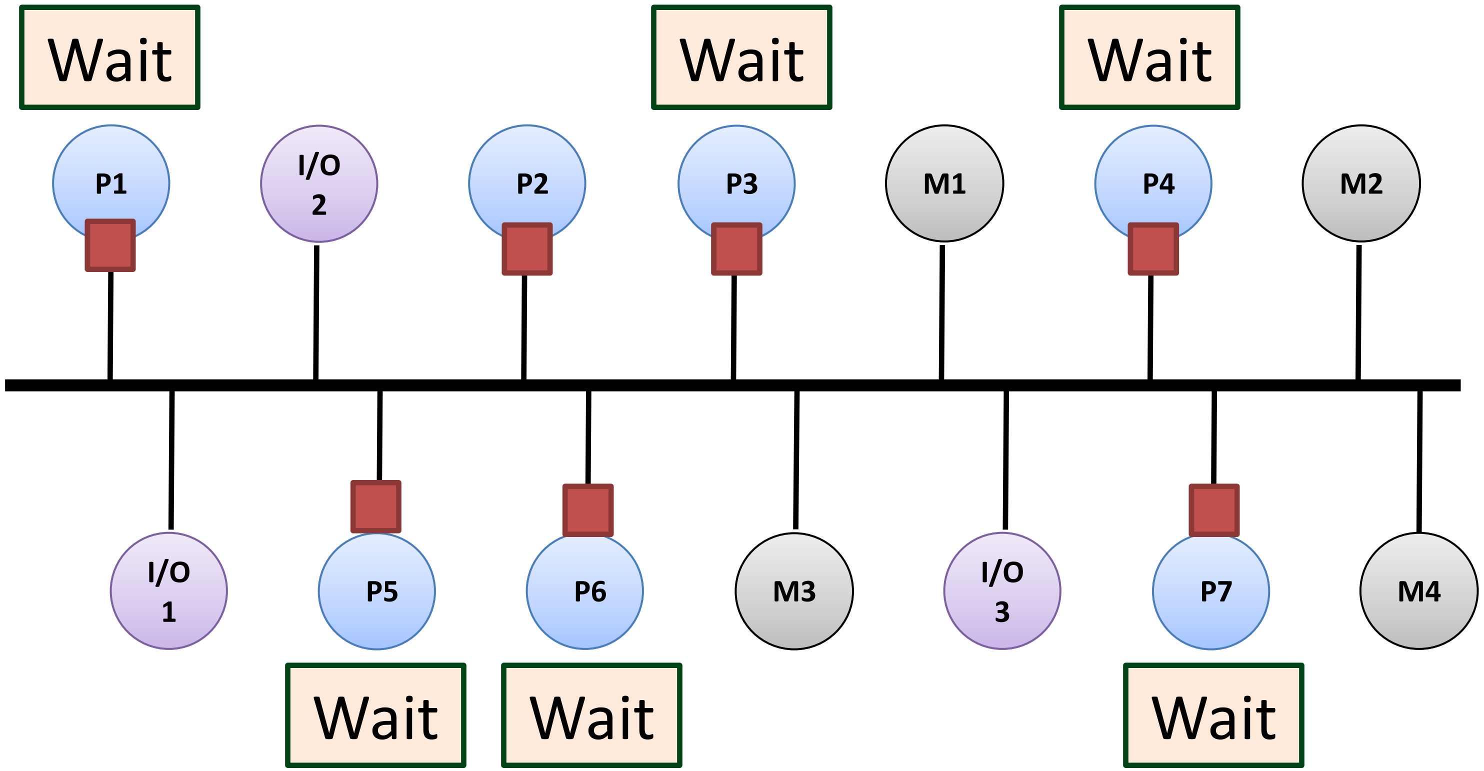

Hierarchical Bus Communication

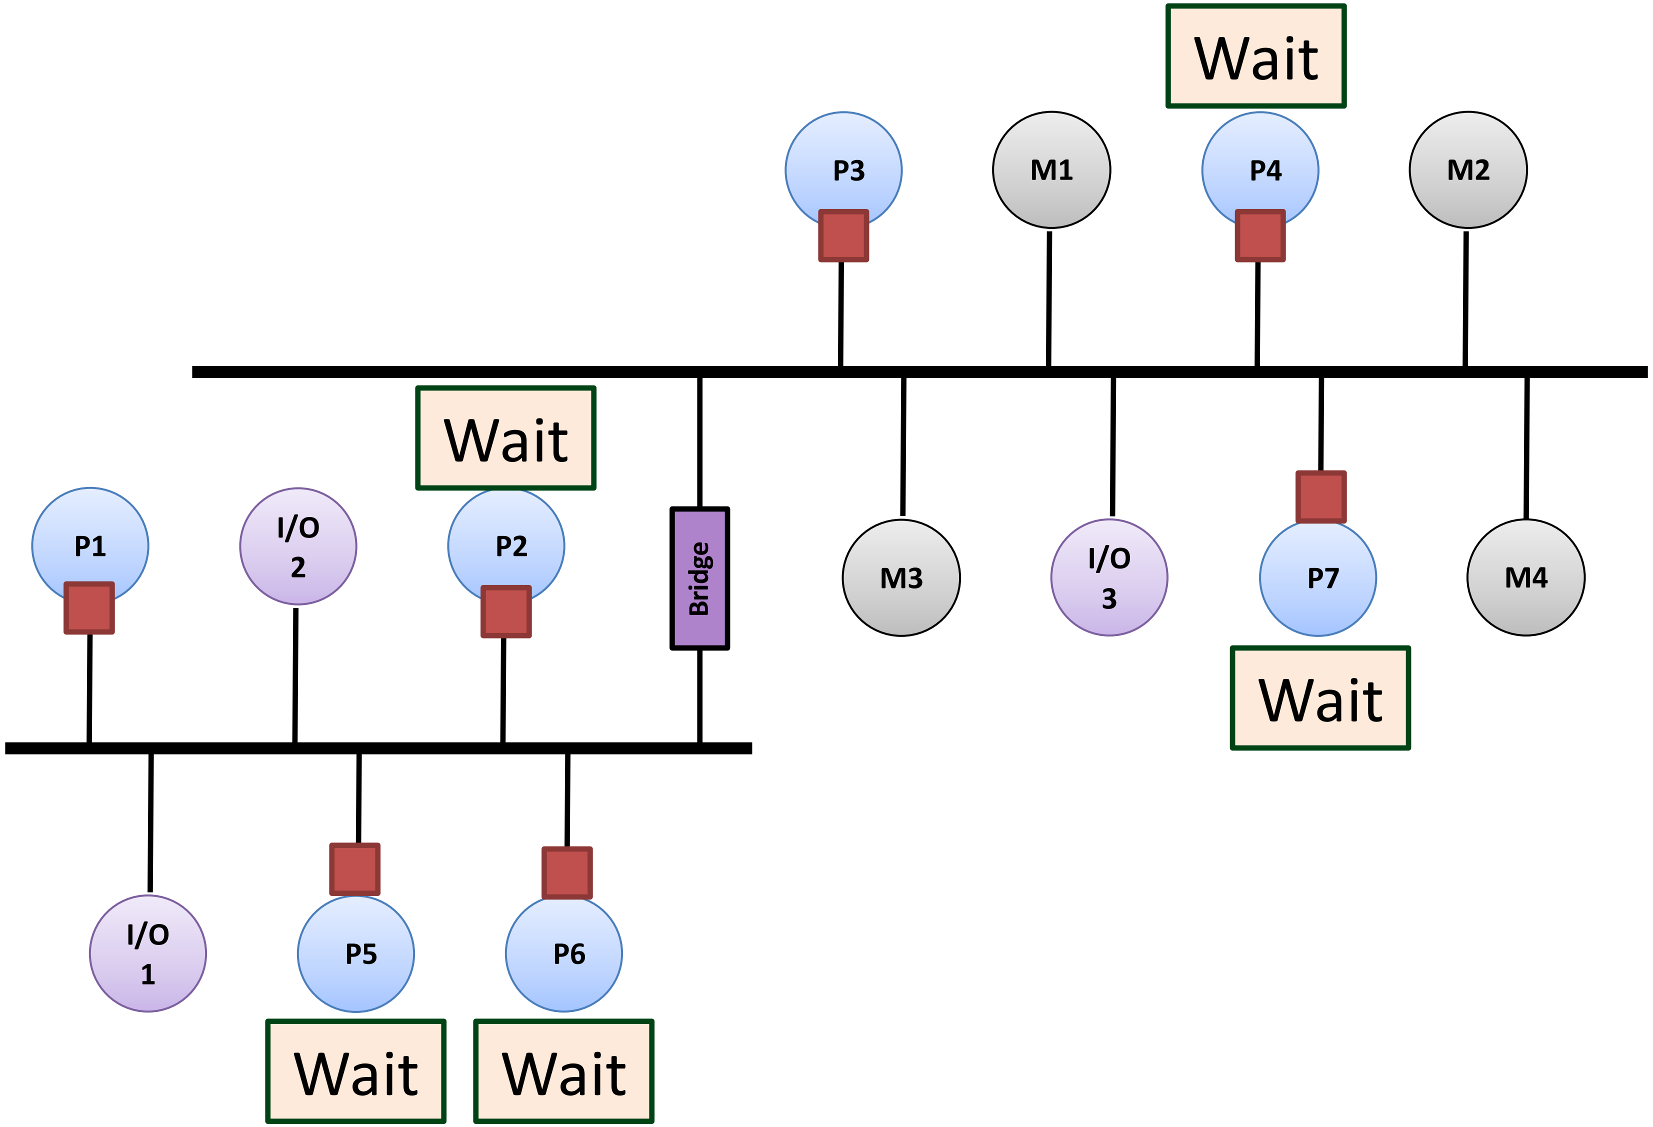

Common Bus Communication

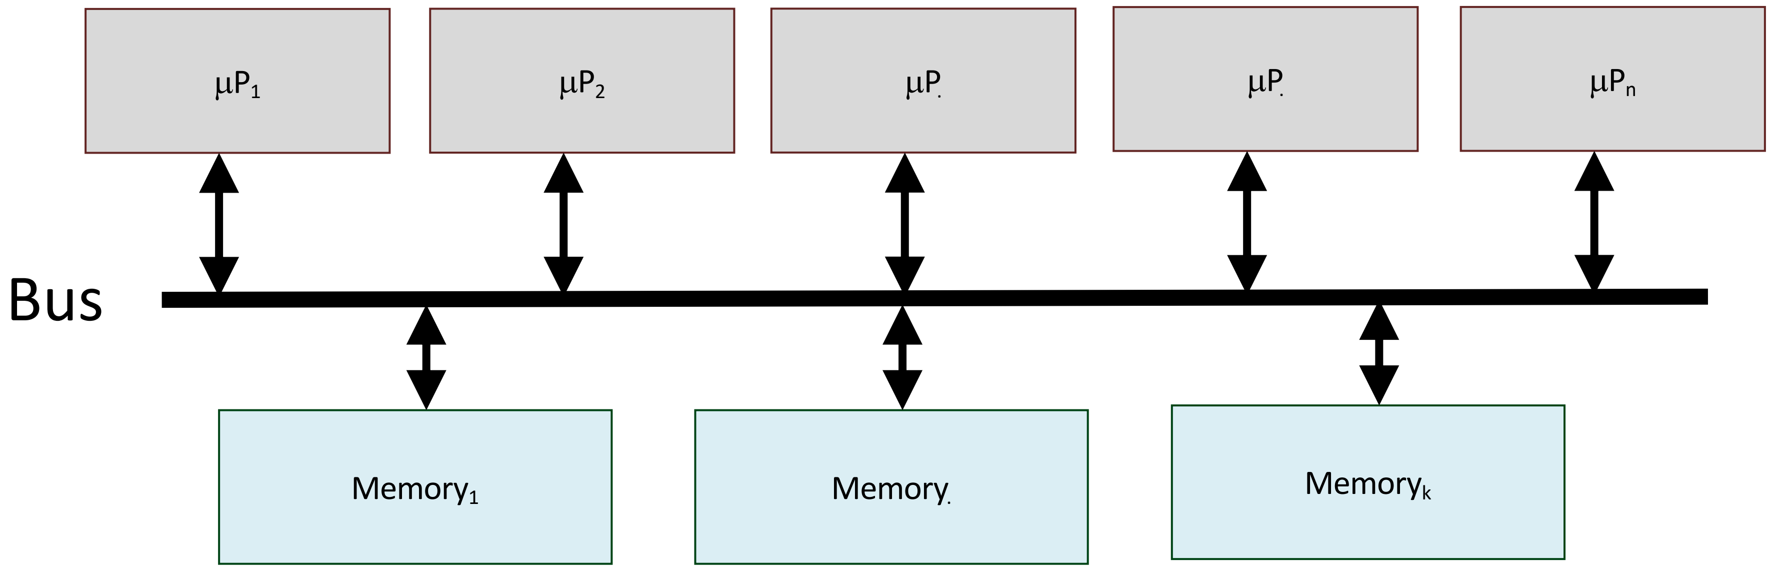

- Advantages
  - The concepts are simple and well understood
  - Any bus is almost directly compatible with most available IPs

Common Bus Communication

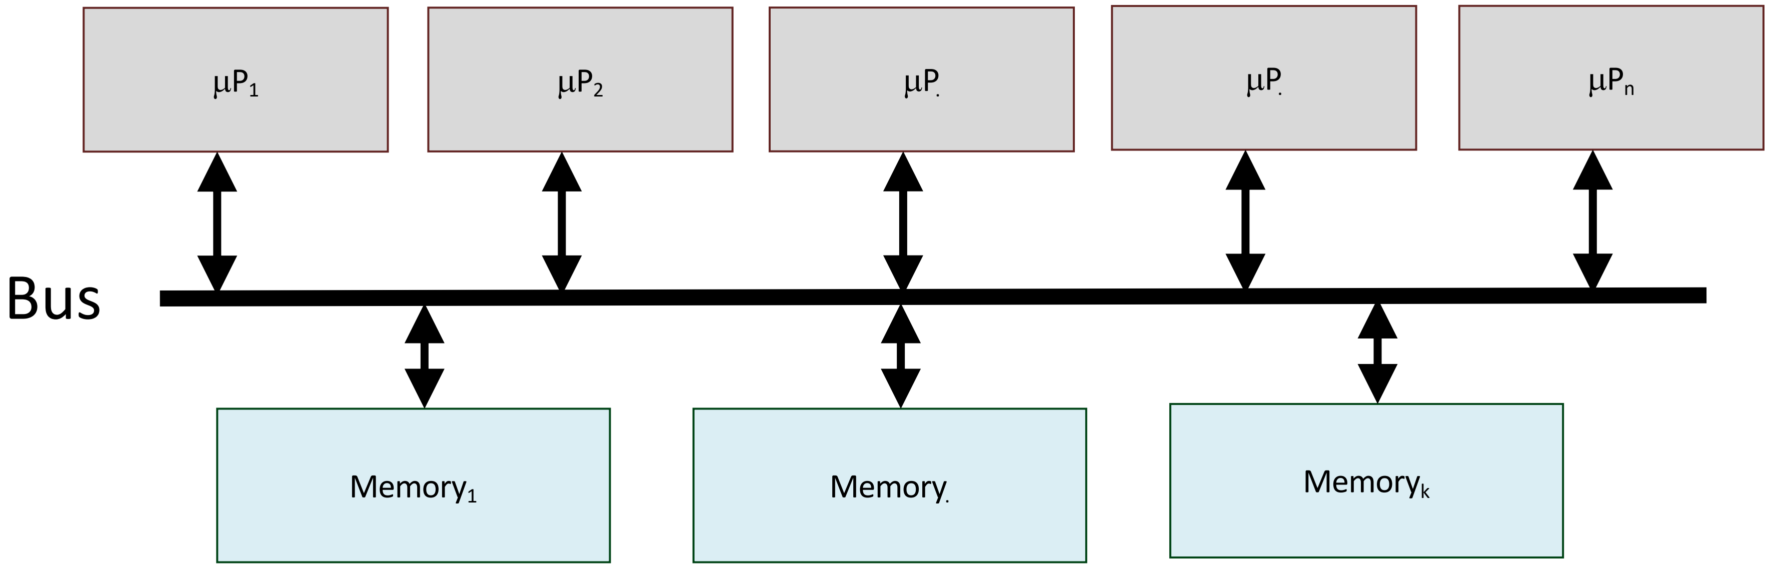

- Disadvantages
  - Scalability issues: does not scale beyond 8-16 cores
  - Electrical loading on the bus significantly reduces its speed
    - Bus arbiter delay grows with the number of processors
  - The shared bus cannot support the bandwidth demand

Matrix Bus Communication

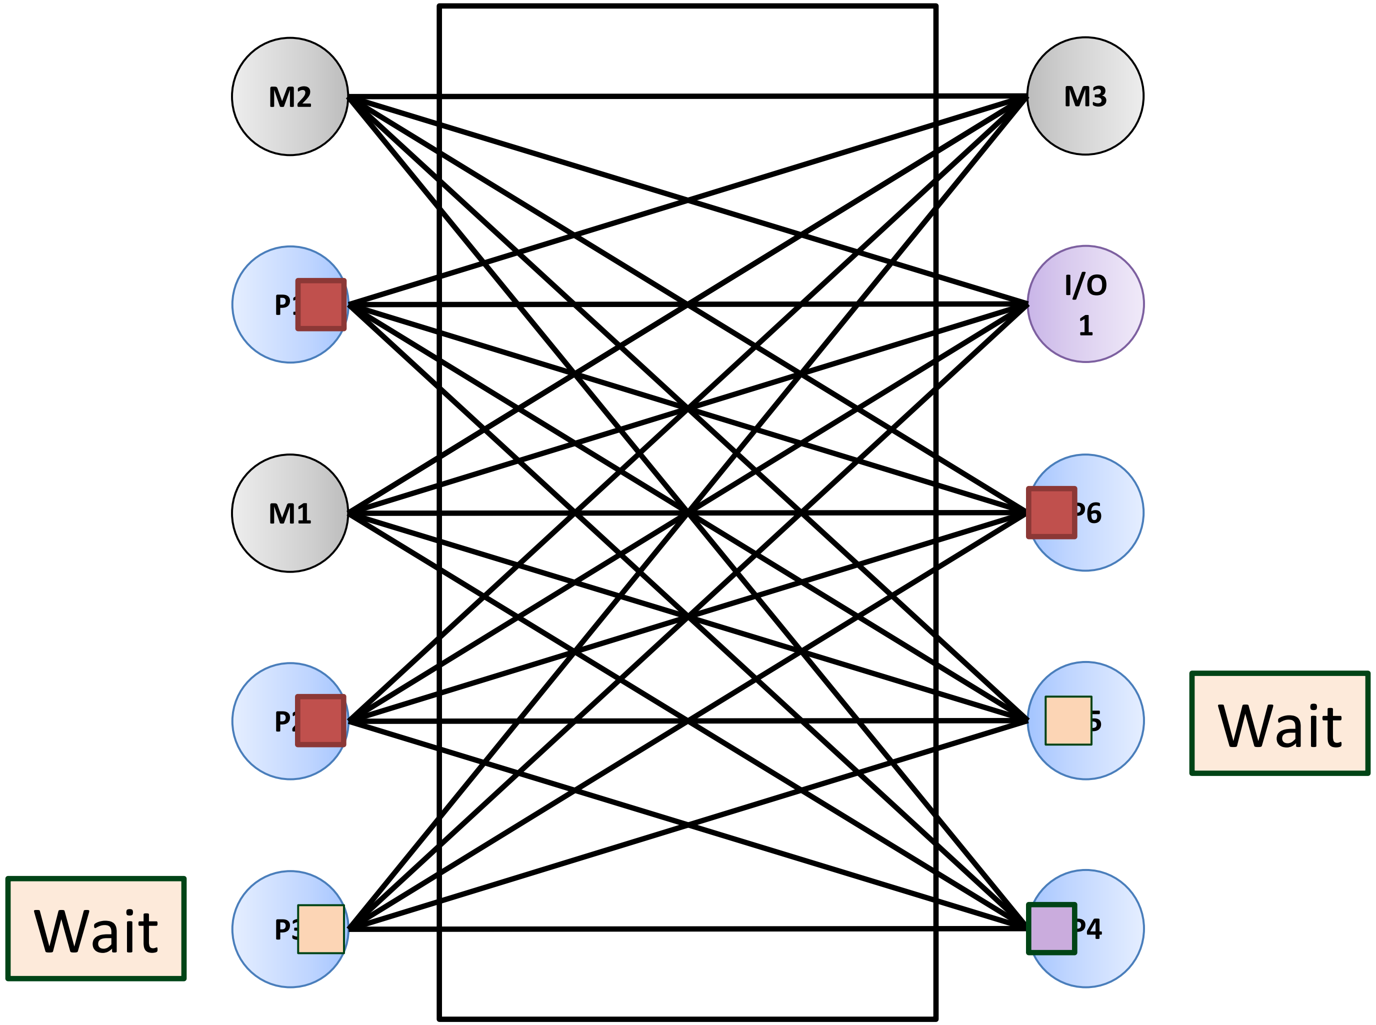

On-Chip Networks

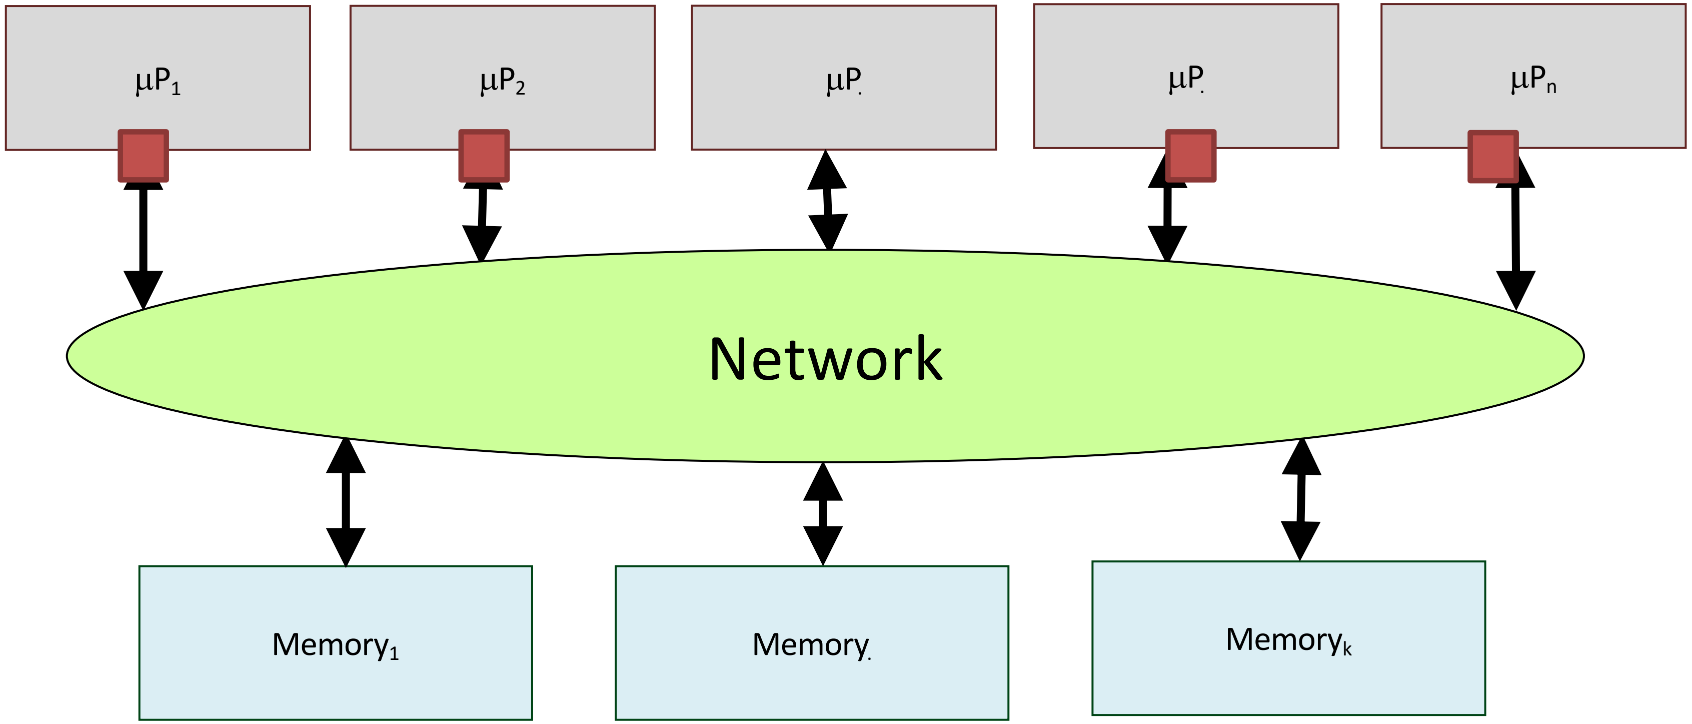

- Higher bandwidth with more concurrent communications
- More scalable with better electrical properties

Architecture of Interconnection Networks

- How to connect the nodes up (processors, memories, router line cards, SoC modules)?
  - Topology
- Which path should a message take?
  - Routing and deadlock
- How is the message actually forwarded from source to destination?
  - Flow Control

Architecture of Interconnection Networks

- How to build the routers?
  - Router microarchitecture
- How to build the links?
  - Link Architecture
- How do processing core talk to the network?
  - Network Interface

Network-on-Chip (NoC)

- An NoC or OIN (On-chip Interconnect Network)
- An NoC architecture is defined by
  - Its topology
    - The physical organization of nodes in the network
  - Its flow control mechanism
    - Which establishes the data formatting, the switching protocol and the buffer allocation
  - Its routing algorithm
    - Which determines the path selected by a packet to reach its destination under a given application

Architecture of Interconnection Networks

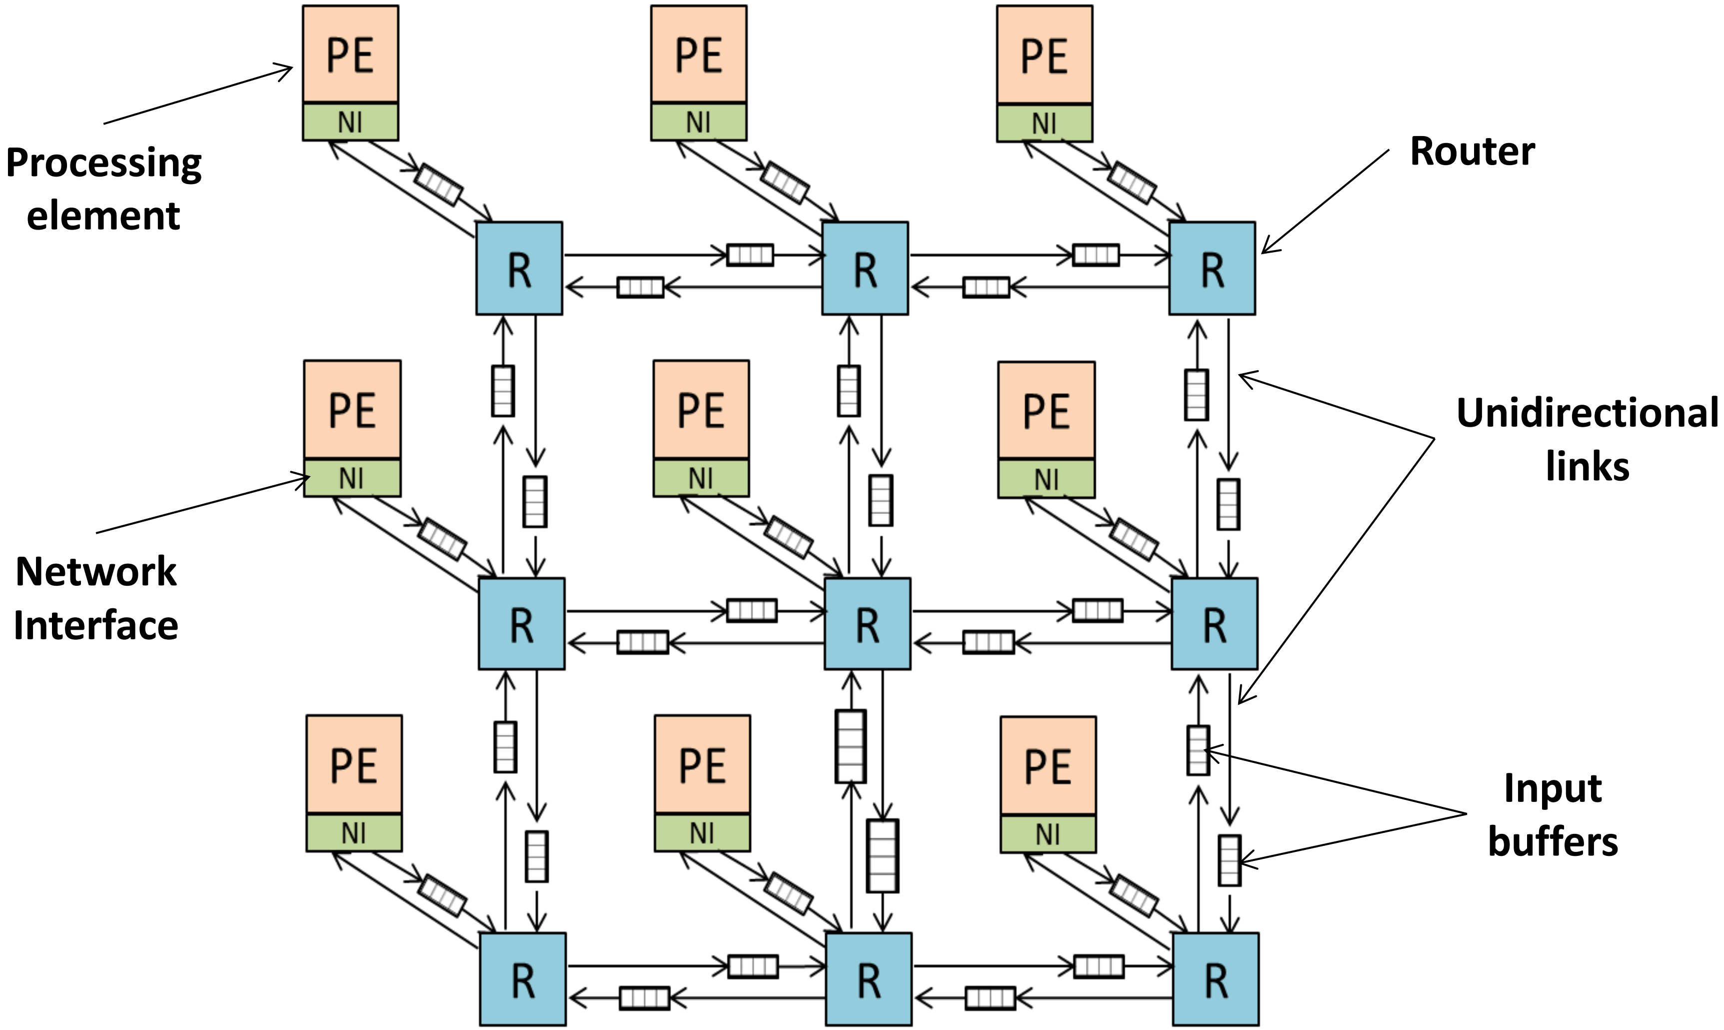

Topological Properties

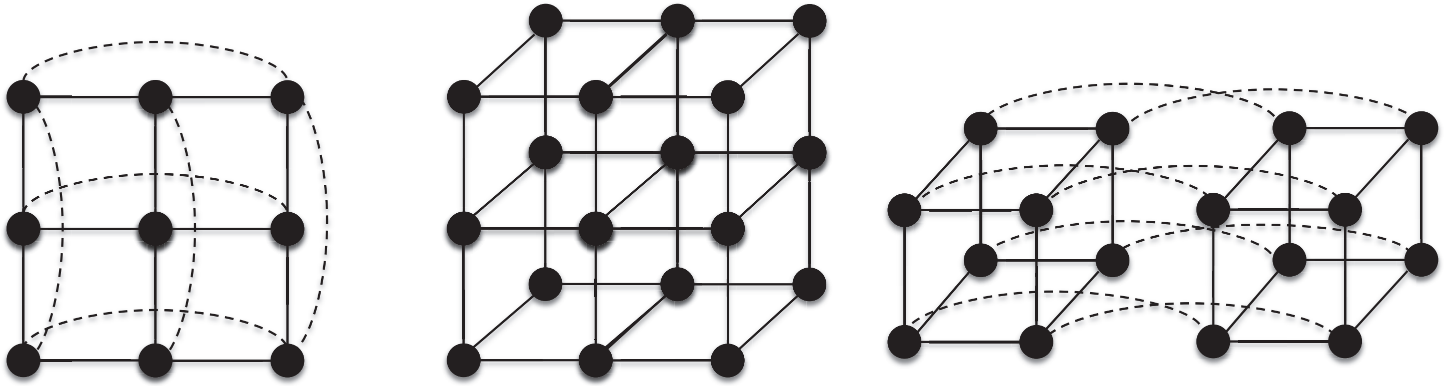

- Routing Distance: number of links on route
- Diameter: maximum routing distance
- Average Distance
- A network is partitioned by a set of links if their removal disconnects the graph
- Bisection Bandwidth: is the bandwidth crossing a minimal cut that divides the network in half

More Topology Examples

- 2D mesh is most popular topology

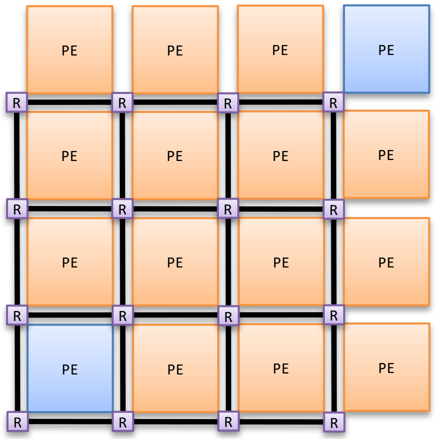

- k-ary n-fly butterfly network
  - Blocking multi-stage network

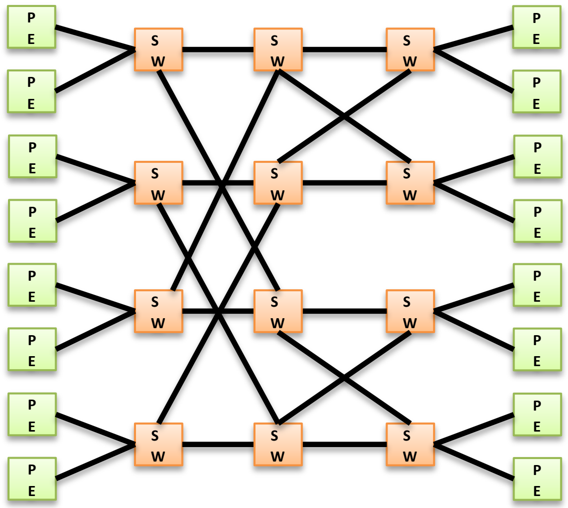

More Topology Examples

- 2D mesh is most popular topology

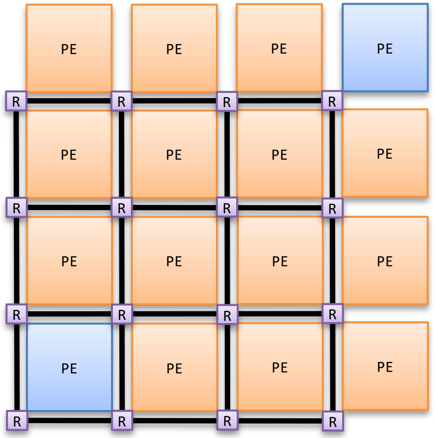

- k-ary n-fly butterfly network
  - Blocking multi-stage network
- Fat tree topology

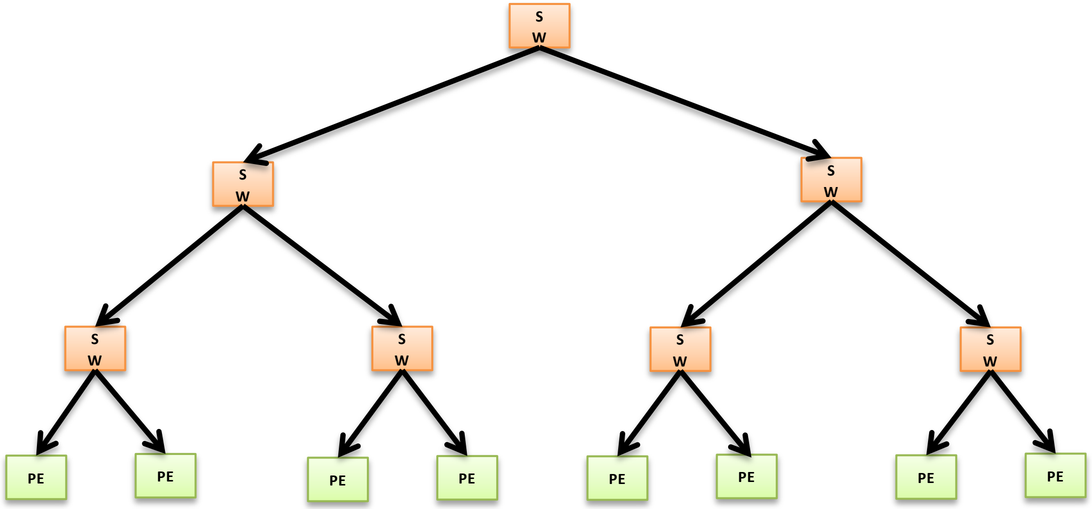

Flow Control

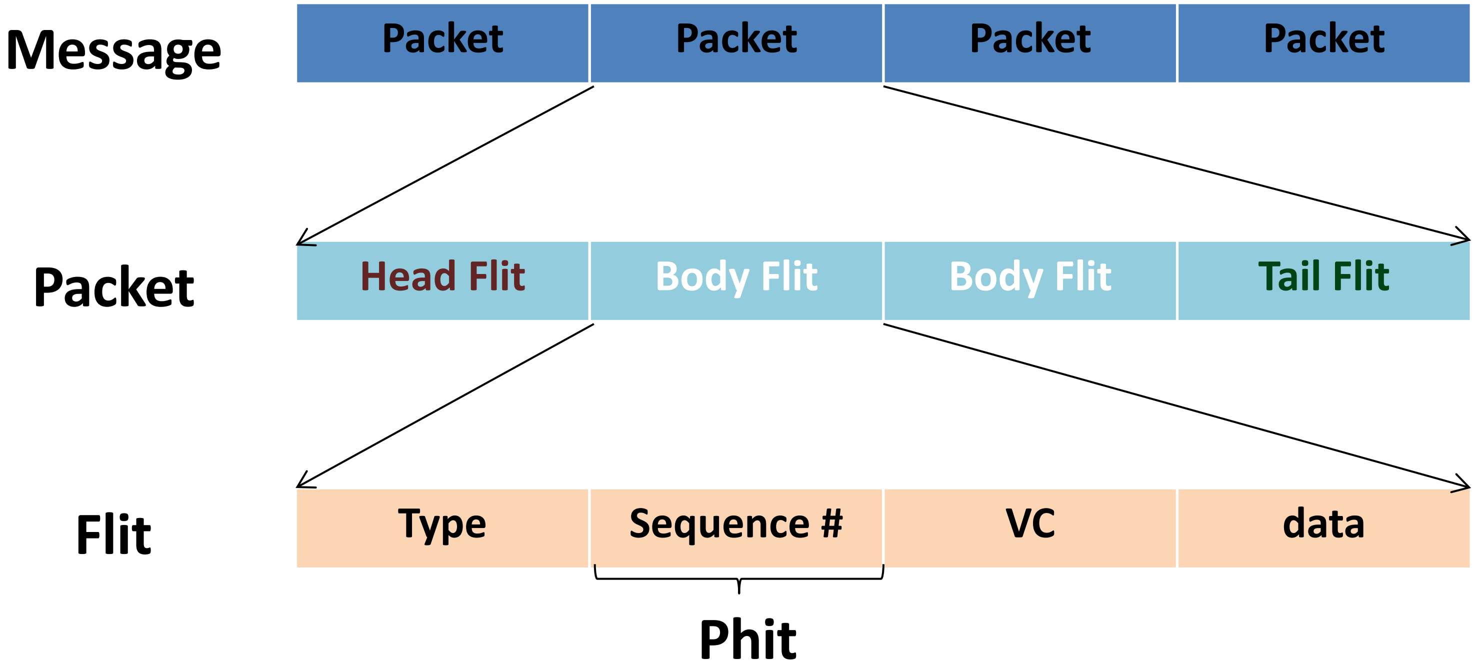

- Flit: flow control digit (basic unit of bandwidth/storage allocation)
- Phit: physical transfer digit (transferred in single clock)

Flow control protocols

- Bufferless
  - Dropping
  - Misrouting
  - Circuit switching
- Buffered
  - Store-and-forward
  - Virtual cut-through
  - Wormhole

Complexity & Efficiency

Resource Contention

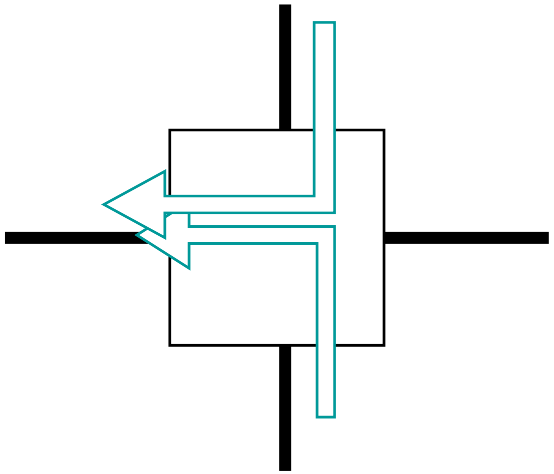

- Two packets trying to use the same link at the same time
  - Limited or no buffering
  - Drop?
- Problem arises because packages are sharing resources
  - Bandwidth and buffering

Simplest Flow Control: Dropping

- If two packages arrive an router does not have resources then its drops one
- Flow control protocol on the Internet
- Not used in interconnection networks – why?

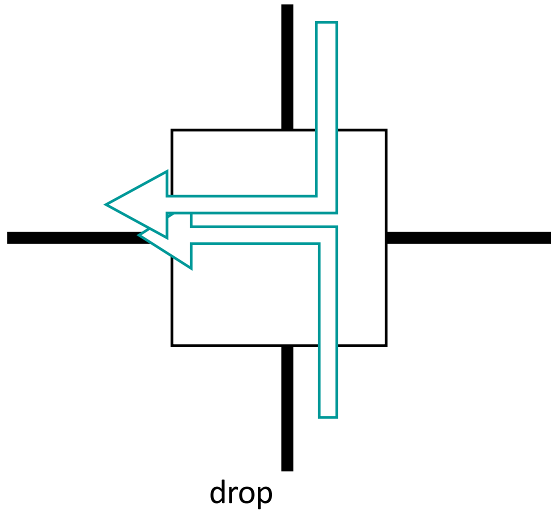

Simplest Flow Control: Misrouting

- Philosophy behind misrouting: intentionally route away from congestion
- No need for buffering
- Problems?
  - Livelock: need to guarantee that progress is made

Circuit Switching

- Characteristics
  - Bufferless
  - Probe that sets up path through network
  - Reserve all links
  - Data sent through links
  - Form a circuit from source to destination
- Advantage: low latency transfers, once path is reserved
- Disadvantage: pure circuit switching does not scale well with NoC size

Virtual Channel Router

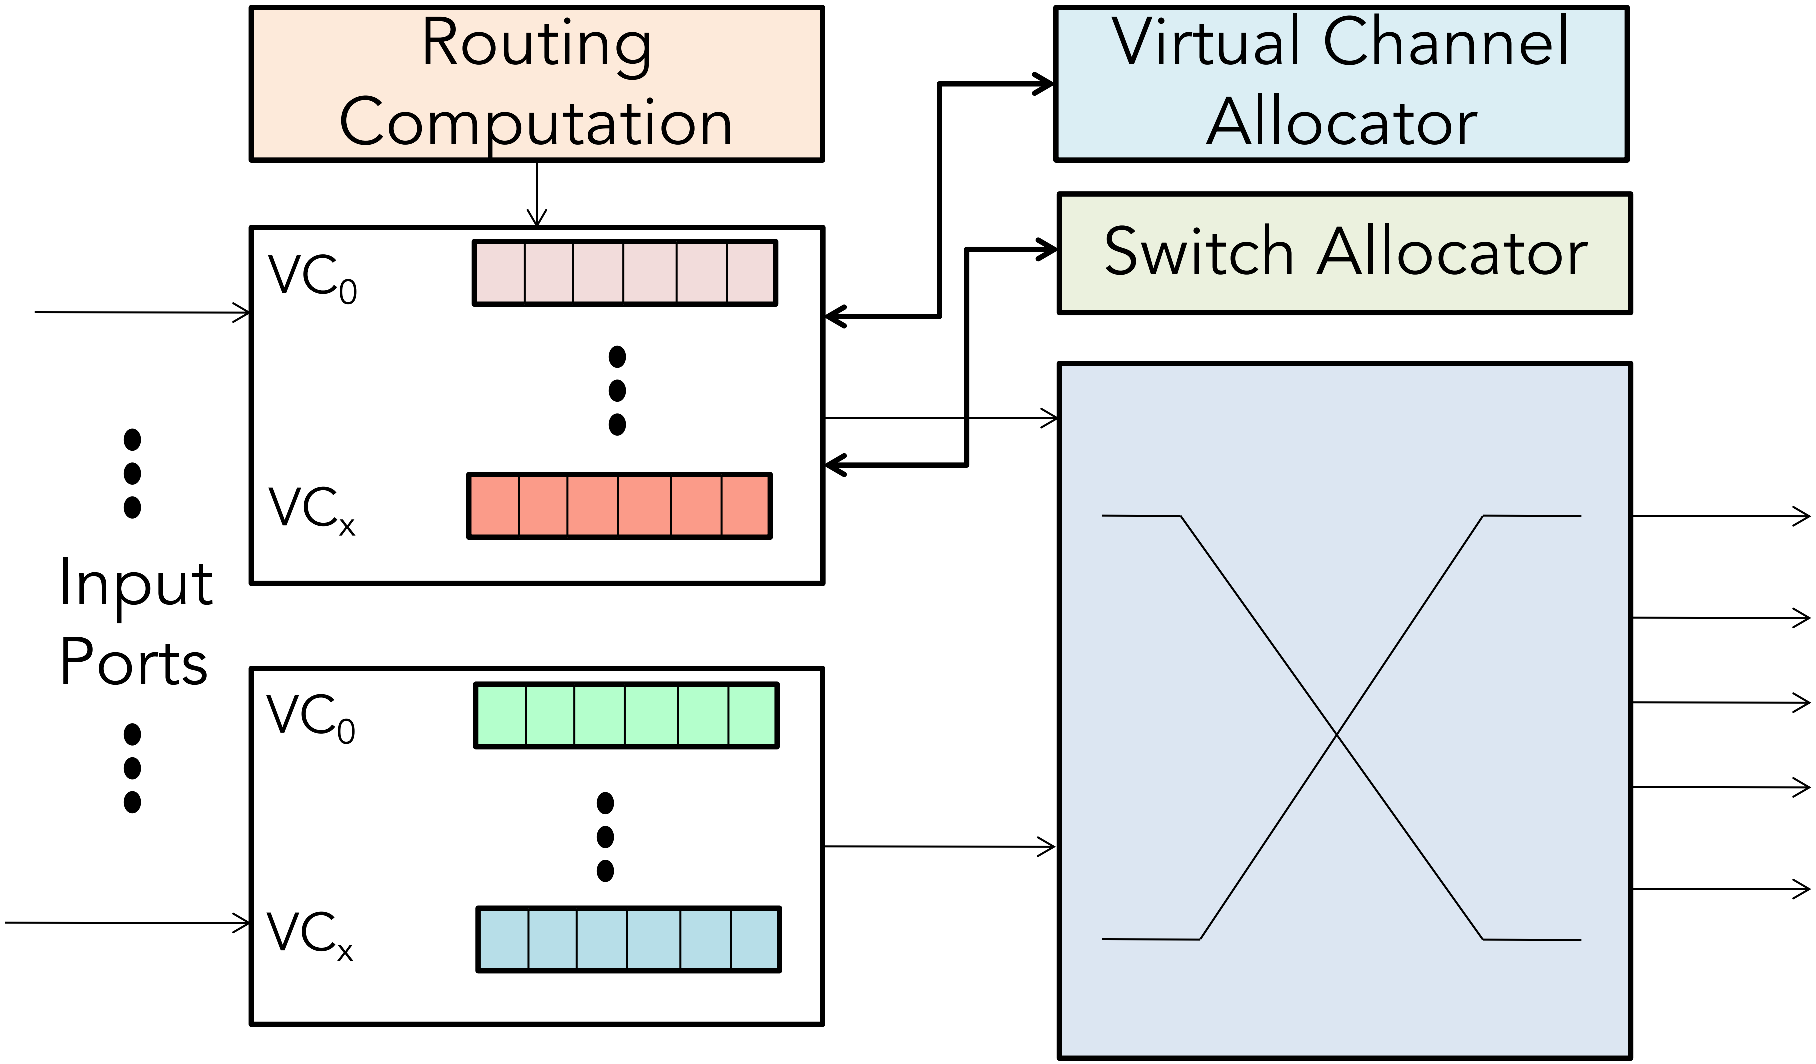

Virtual Channel Router

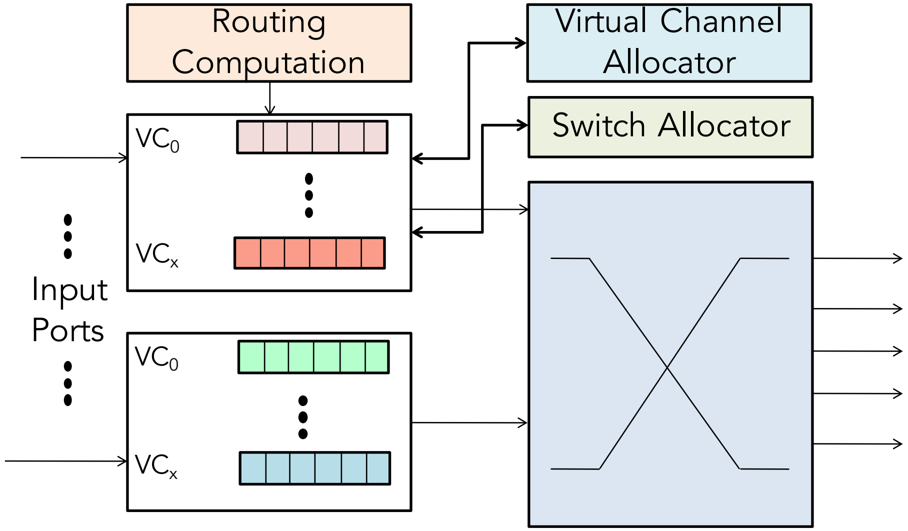

- The routing operation takes four steps or phases
  - Route computation (RC)
    - When a head flit (the first flit of a packet) arrives at an input channel, the router stores the flit in the buffer for the allocated virtual channel and determines the next hop for the packet

Virtual Channel Router

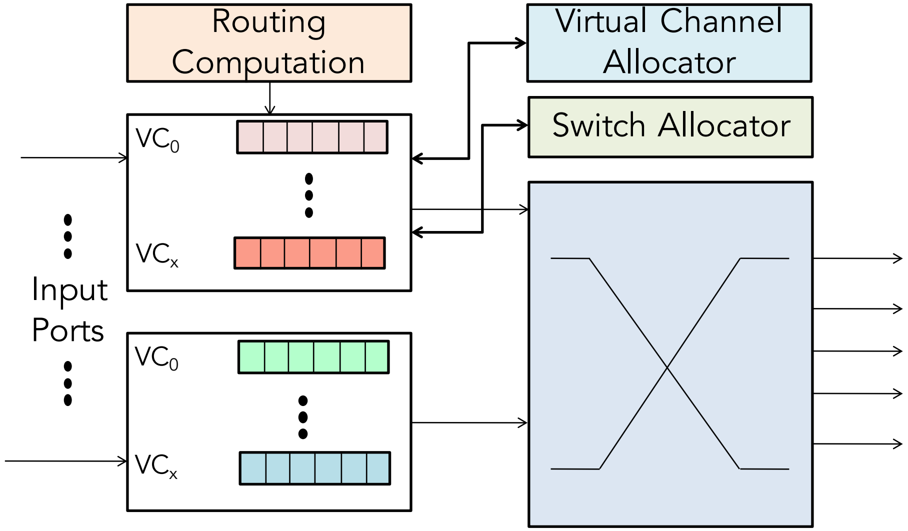

- The routing operation takes four steps or phases
  - Route computation (RC)
    - When a head flit (the first flit of a packet) arrives at an input channel, the router stores the flit in the buffer for the allocated virtual channel and determines the next hop for the packet
  - Virtual-channel allocation (VA)
    - Given the next hop, the router then allocates a virtual channel in the next hop

Virtual Channel Router

- The routing operation takes four steps or phases
  - Route computation (RC)
    - When a head flit (the first flit of a packet) arrives at an input channel, the router stores the flit in the buffer for the allocated virtual channel and determines the next hop for the packet
  - Virtual-channel allocation (VA)
    - Given the next hop, the router then allocates a virtual channel in the next hop
  - Switch allocation (SA)
    - Finally, the flit competes for a switch
  - Switch traversal (ST)
    - If the next hop can accept the flit, it is then moved to the output port

Store-and-Forward

- Head flits makes intermediate stops and waits for the whole package to arrive before moving to next hop
- Other messages can use intermediate links
- Buffering allows packet to wait for channel
- Drawback?
  - Serialization latency experienced at each hop/channel

Buffers

0	H	B	B	B	T															
1						H	B	B	B	T										
2											H	B	B	B	T					
3																H	B	B	B	T

Virtual Cut-through

- Why wait till entire message has arrived at each intermediate stop?
- The head of the message can dash off first
- When the head gets blocked, whole message gets blocked at one intermediate node
- Used in Alpha 21364

Buffers

0	H	B	B	B	T								
1		H	B	B	B	T							
2			|---- Not Ready ----|					H	B	B	B	T	
3									H	B	B	B	T

Wormhole

- When a message blocks, just block wherever the pieces (flits) of the message are at that time
- Operates like cut-through but with channel and buffers allocated to flits rather than packets
- Channel state (virtual channel) allocated to packet so body flits can follow head flit

Next Learning Module

- On-chip Networking