EECS 570

Lecture 22

Interconnects:

Router Microarchitecture

Fall 2025

Prof. Satish Narayanasamy

http://www.eecs.umich.edu/courses/eecs570/

Slides developed in part by Profs. Adve, Falsafi, Hill, Lebeck, Martin, Narayanasamy, Nowatzyk, Reinhardt, Singh, Smith, Torrellas and Wenisch.
Special acknowledgement to Prof. Jerger of U. Toronto.

Buffer Backpressure

Buffer Backpressure

- Need mechanism to prevent buffer **overflow**
  - Avoid dropping packets
  - Upstream nodes need to know buffer availability at downstream routers
- Significant impact on throughput achieved by flow control
- Two common mechanisms
  - Credits
  - On-off
- Credit-based generally works better
  - On-chip, wires are cheaper than buffers

Credit-Based Flow Control

- Upstream router stores credit counts for each downstream VC

- Upstream router
  - When flit forwarded
    - Decrement credit count
  - Count == 0, buffer full, stop sending

- Downstream router
  - When flit forwarded and buffer freed
    - Send credit to upstream router
    - Upstream increments credit count

Credit Timeline

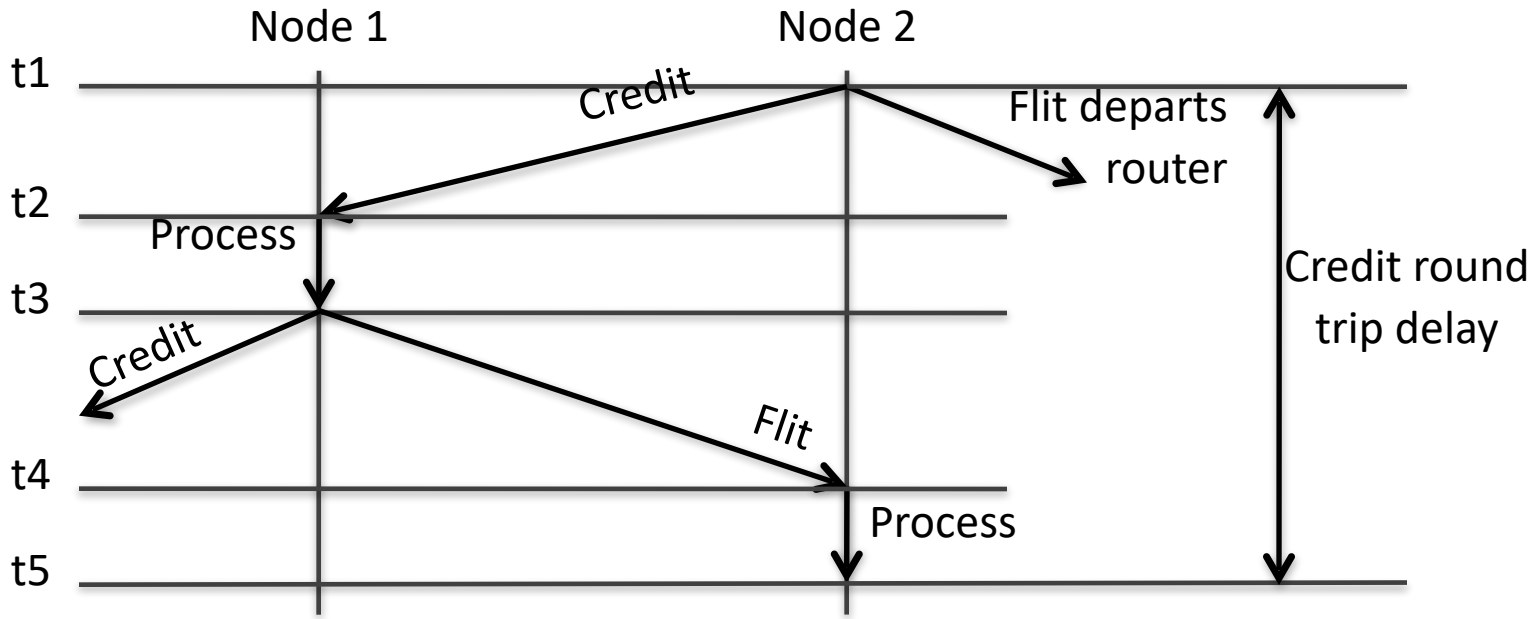

- Round-trip credit delay:
  - Time between when buffer empties and when next flit can be sent from that buffer entry
  - If only single entry buffer, would result in significant throughput degradation
  - Important to size buffers to tolerate credit turn-around

Buffer Utilization

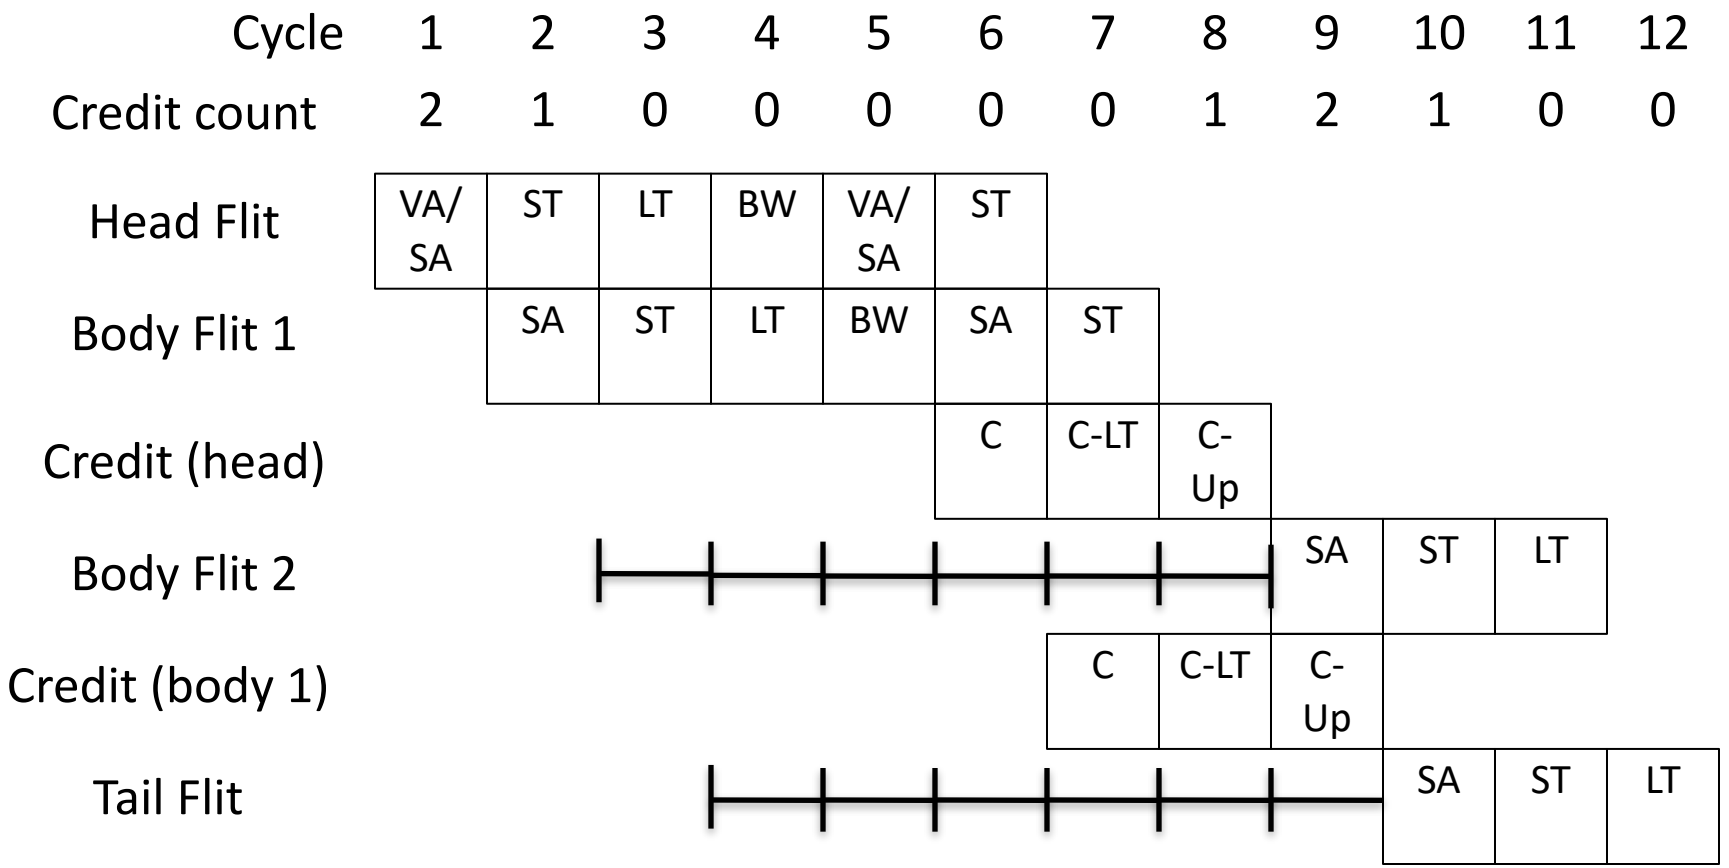

Buffer Sizing

- Prevent backpressure from limiting throughput
  - Buffers must hold # of flits >= turnaround time

- Assume:
  - 1 cycle propagation delay for data and credits
  - 1 cycle credit processing delay
  - 3 cycle router pipeline

- At least 6 flit buffers

On-Off Flow Control

- Credit: requires upstream signaling for every flit
- On-off: decreases upstream signaling
- Off signal
  - Sent when number of free buffers falls below threshold F_{off}
- On signal
  - Sent when number of free buffers rises above threshold F_{on}

On-Off Timeline

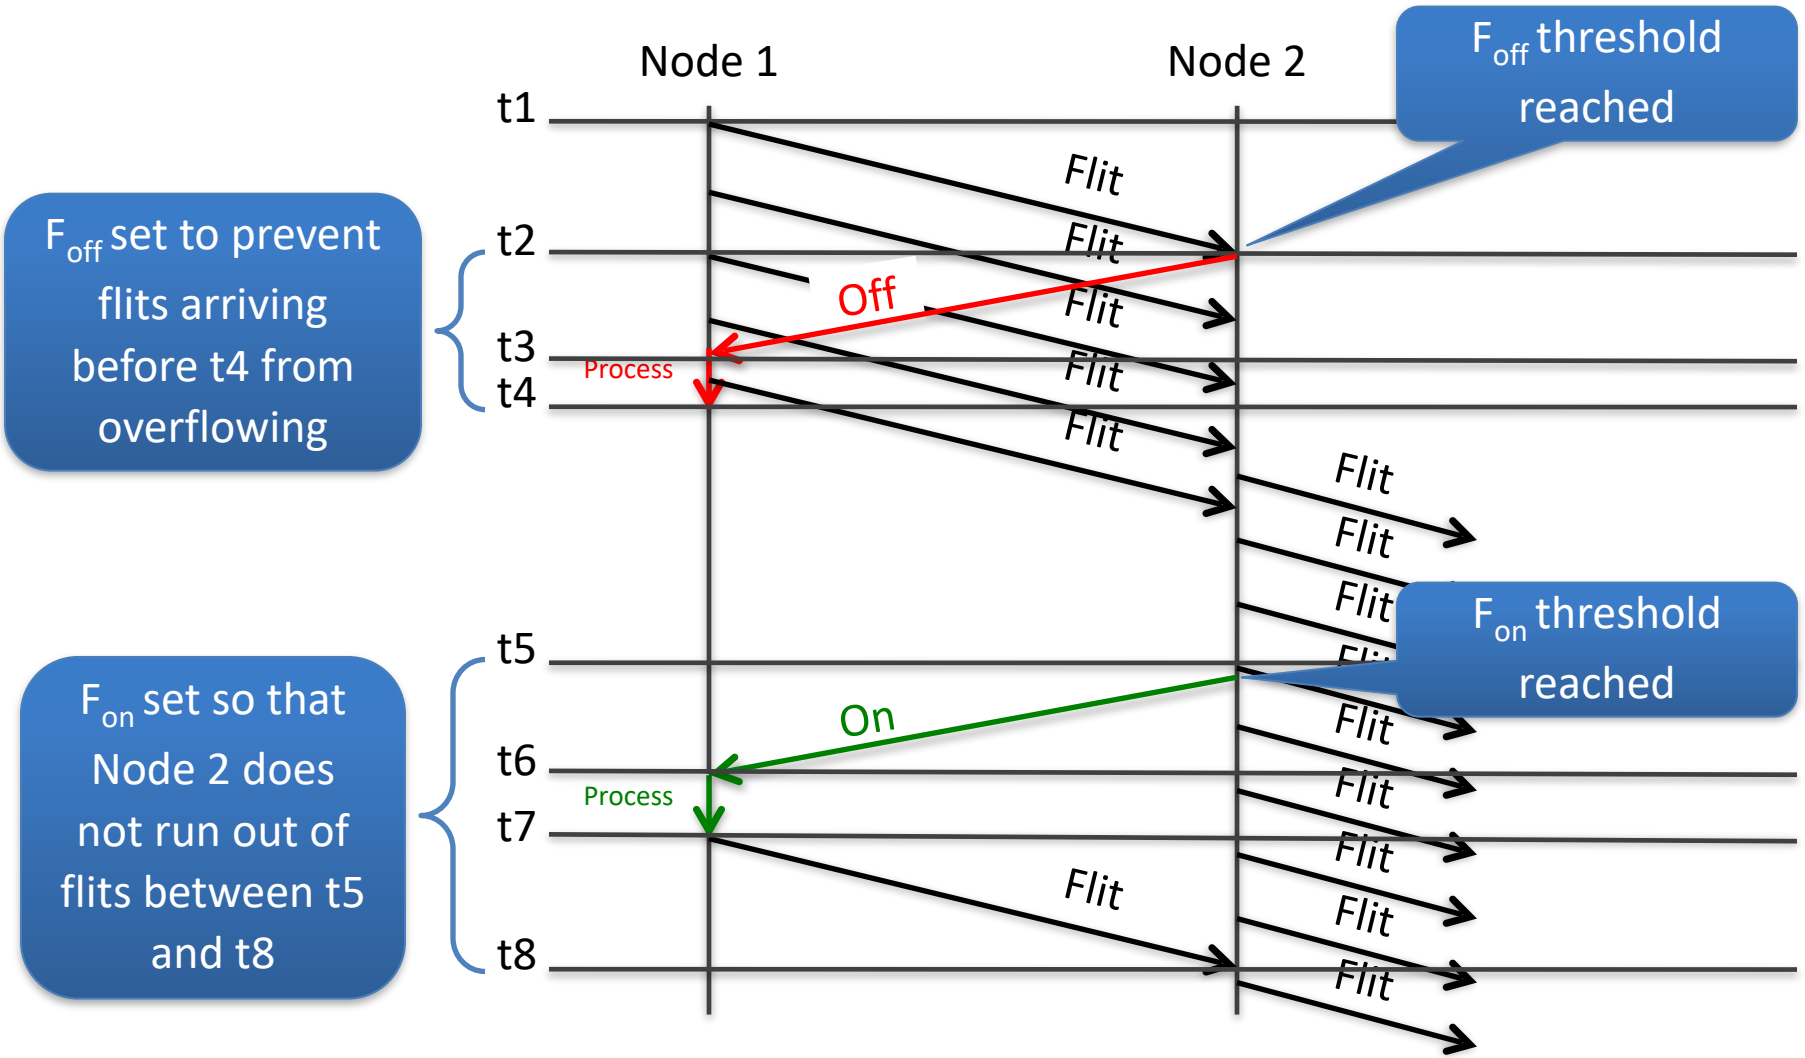

- Less signaling but more buffering
  - On-chip buffers more expensive than wires

Flow Control Summary

- On-chip networks require techniques with lower buffering requirements
  - Wormhole or Virtual Channel flow control
- Avoid dropping packets in on-chip environment
  - Requires buffer backpressure mechanism
- Complexity of flow control impacts router microarchitecture

Router Microarchitecture

- Topology: connectivity
- Routing: paths
- Flow control: resource allocation

- Router Microarchitecture
  - Implementation of routing, flow control and router pipeline
  - Impacts per-hop delay and energy

Router Microarchitecture Overview

- Focus on microarchitecture of Virtual Channel router
  - Router complexity increase with bandwidth demands
  - Simple routers built when high throughput is not needed
    - Wormhole flow control, unpipelined, limited buffer

Virtual Channel Router

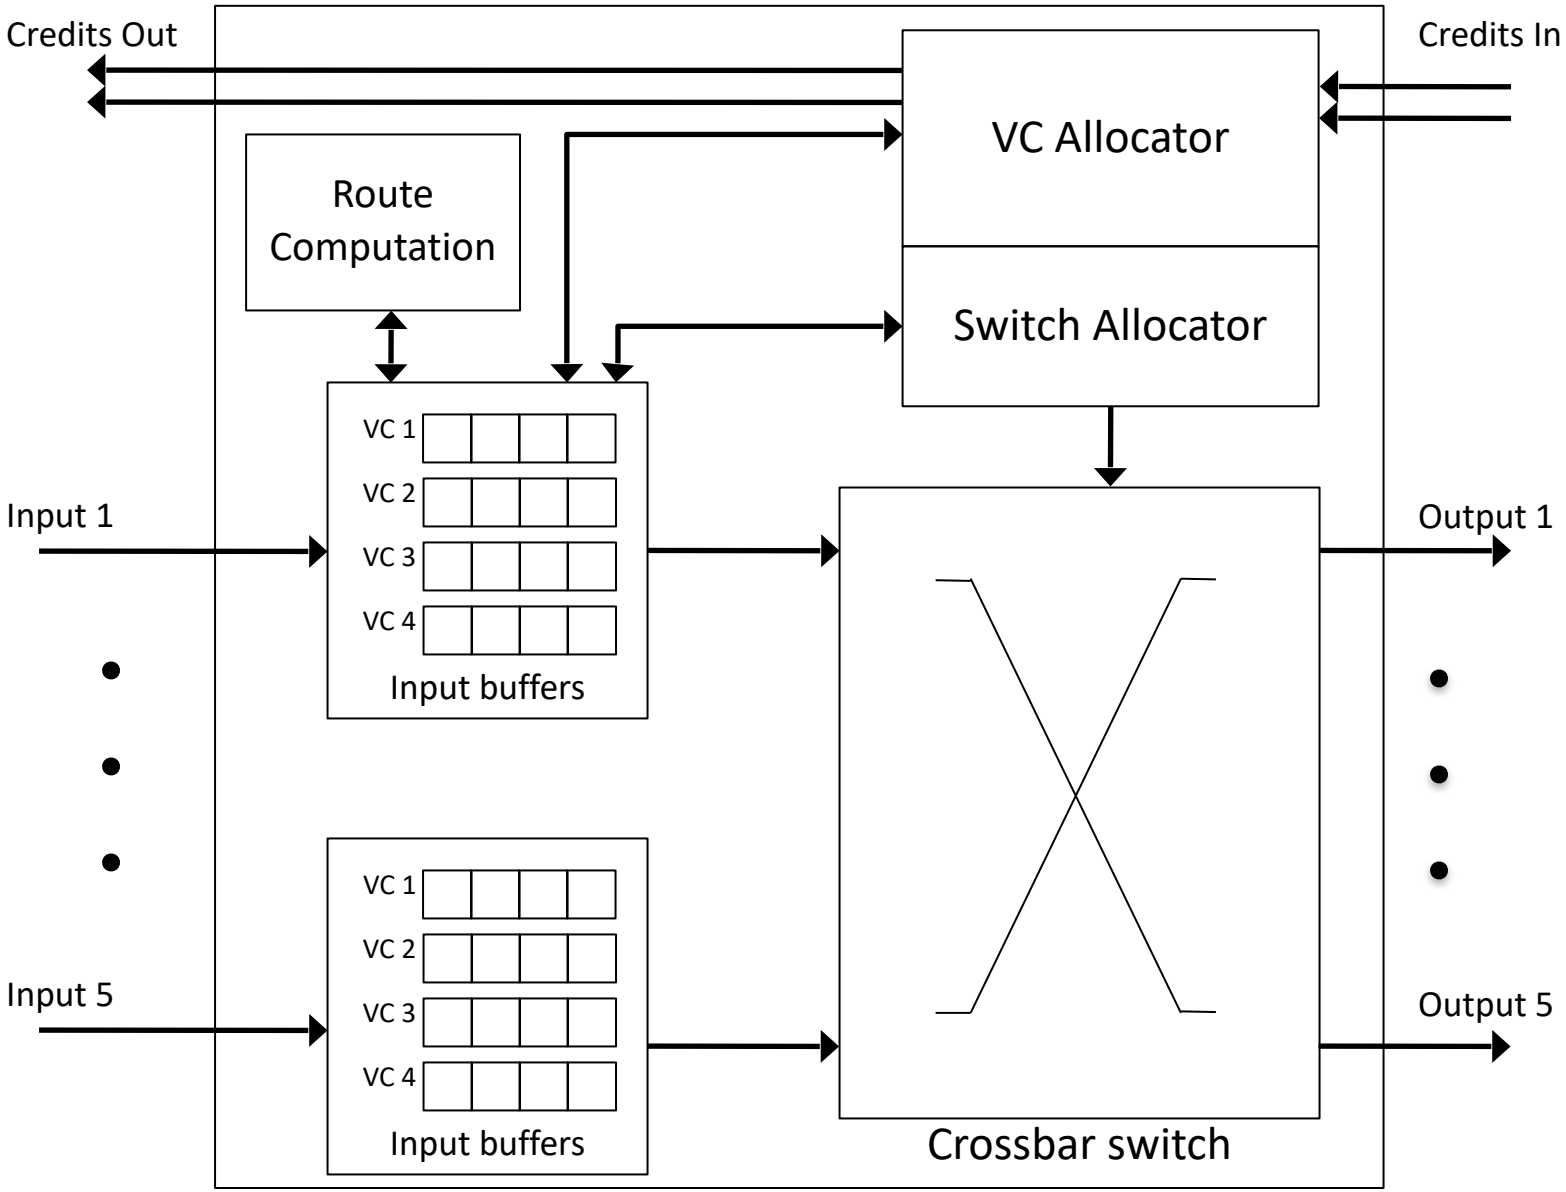

Router Components

- Input buffers, route computation logic, virtual channel allocator, switch allocator, crossbar switch
- Most OCN routers are input buffered
  - Use single-ported memories
- Buffer store flits for duration in router

Baseline Router Pipeline

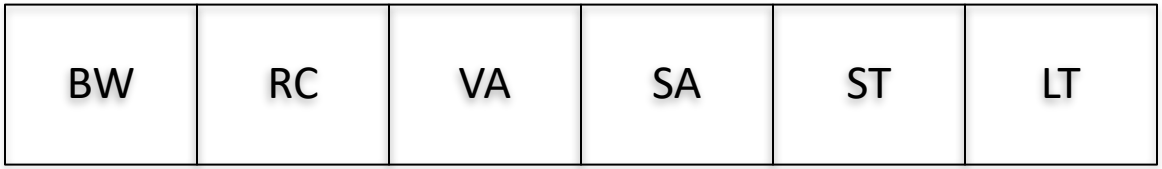

- Logical stages
  - Fit into physical stages depending on frequency
- Canonical 5-stage pipeline
  - BW: Buffer Write
  - RC: Routing computation
  - VA: Virtual Channel Allocation
  - SA: Switch Allocation
  - ST: Switch Traversal
  - LT: Link Traversal

Baseline Router Pipeline (2)

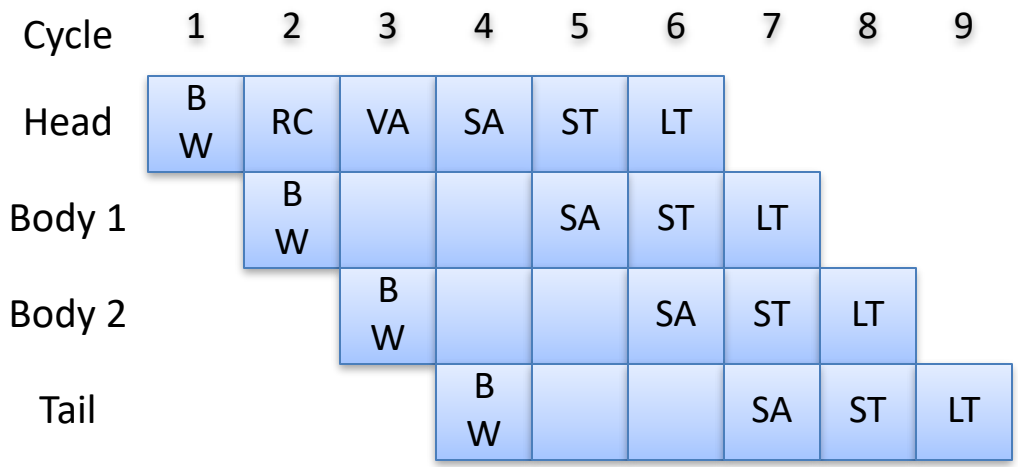

- Routing computation performed once per packet
- Virtual channel allocated once per packet
- Body and tail flits inherit this info from head flit

Atomic Modules and Dependencies in Router

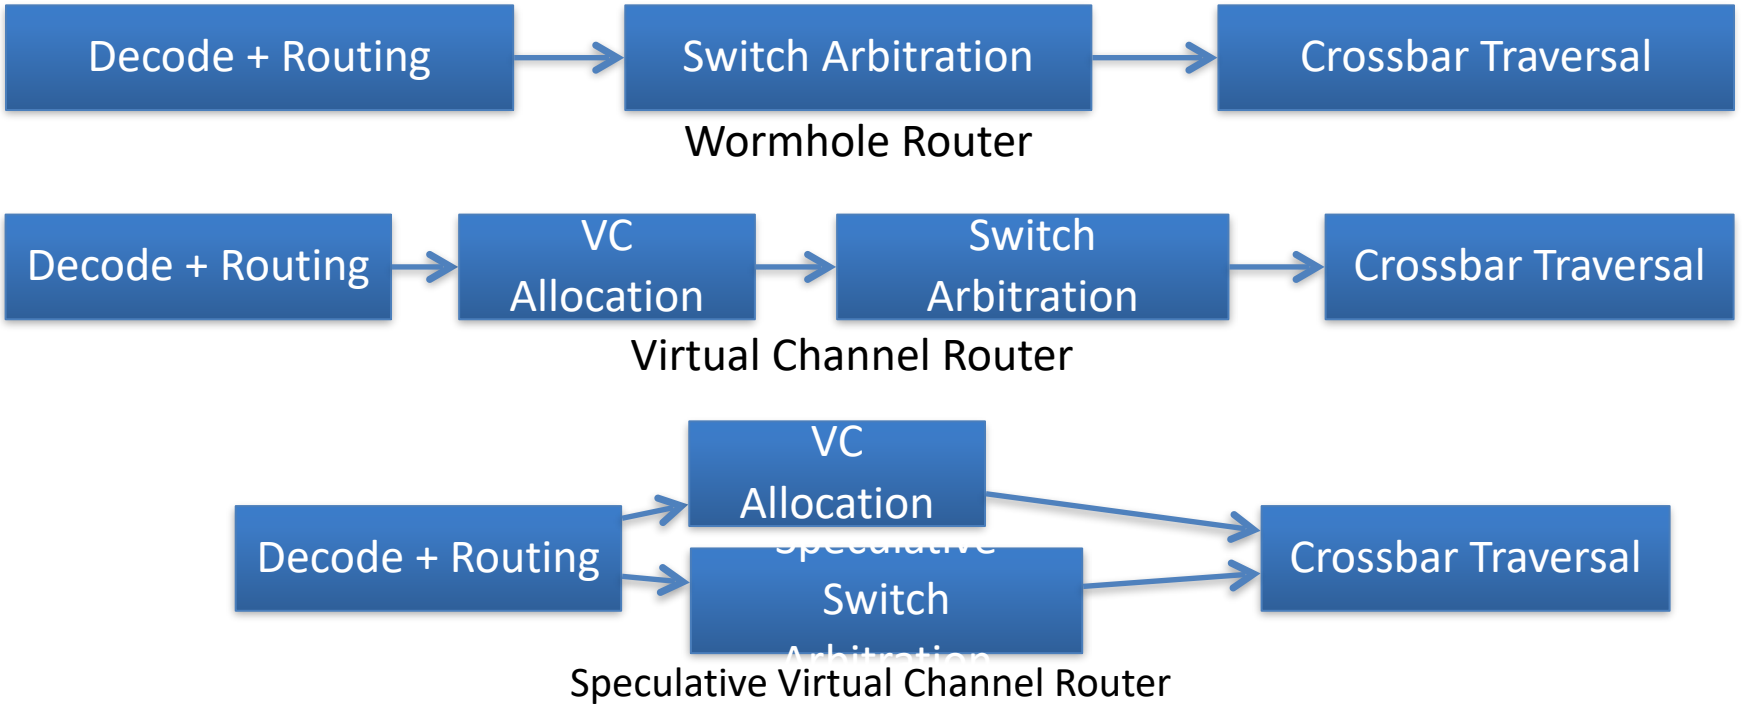

- Dependence between output of one module and input of another
  - Determine critical path through router
  - Cannot bid for switch port until routing performed

Router Pipeline Performance

- Baseline (no load) delay
- Ideally, only pay link delay
- Techniques to reduce pipeline stages
  - Saves latency and energy
  - Shallow pipelines also lower buffer turnaround time

Pipeline Optimizations: Lookahead Routing

- At current router perform route computation for next router
  - Allows your RC to overlap with Buffer Write (BW)
  - Can compute route for next hop in parallel with your VA stage

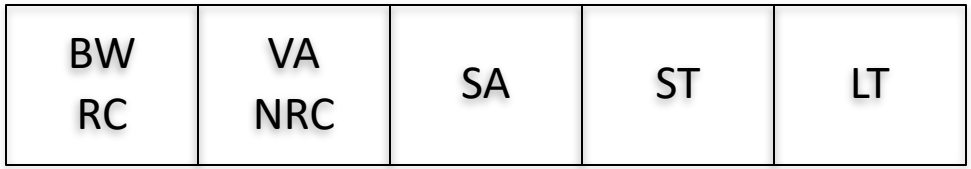

  - Precomputing route allows flits to compete for VCs immediately after BW

Pipeline Optimizations: Speculation

- Assume that Virtual Channel Allocation stage will be successful
  - Valid under low to moderate loads
- Entire VA and SA in parallel

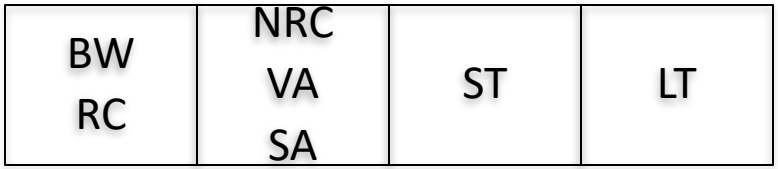

- If VA unsuccessful (no virtual channel returned)
  - Must repeat VA/SA in next cycle
- Prioritize non-speculative requests

Pipeline Optimizations: Bypassing

- When no flits in input buffer
  - Speculatively enter ST
  - On port conflict, speculation aborted
    - Flit written into buffer (BW) and performs SA

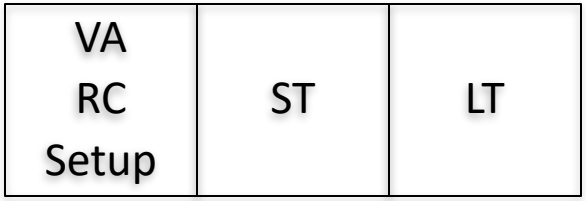

  - In the first stage, a free VC is allocated, next routing is performed and the crossbar is setup

Pipeline Bypassing

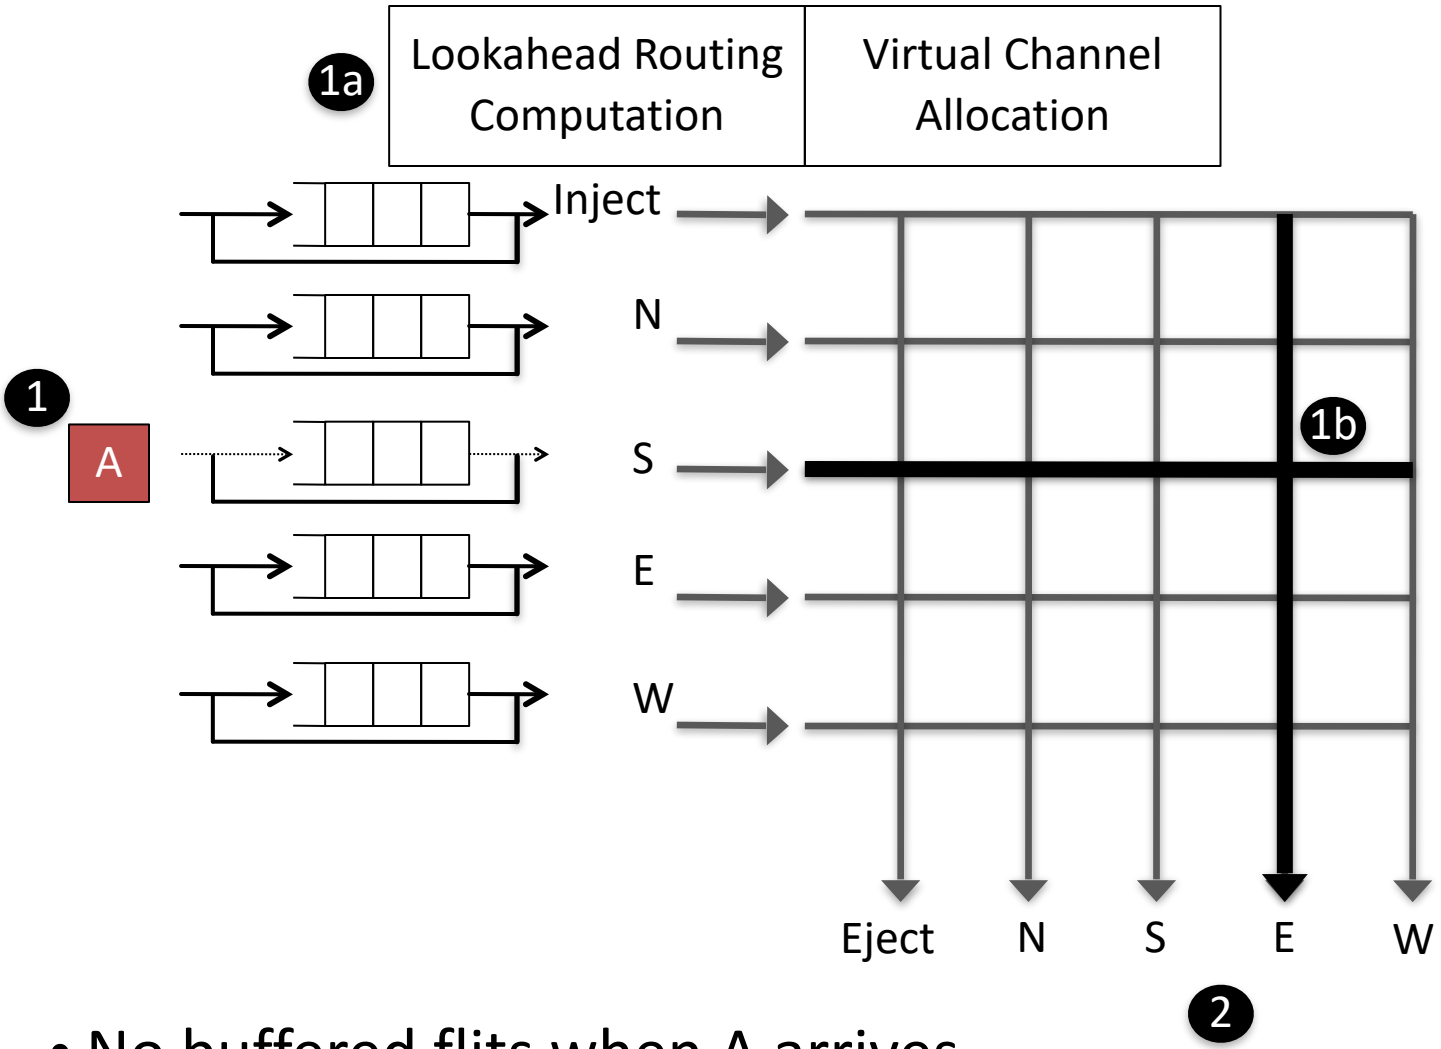

- No buffered flits when A arrives

Speculation

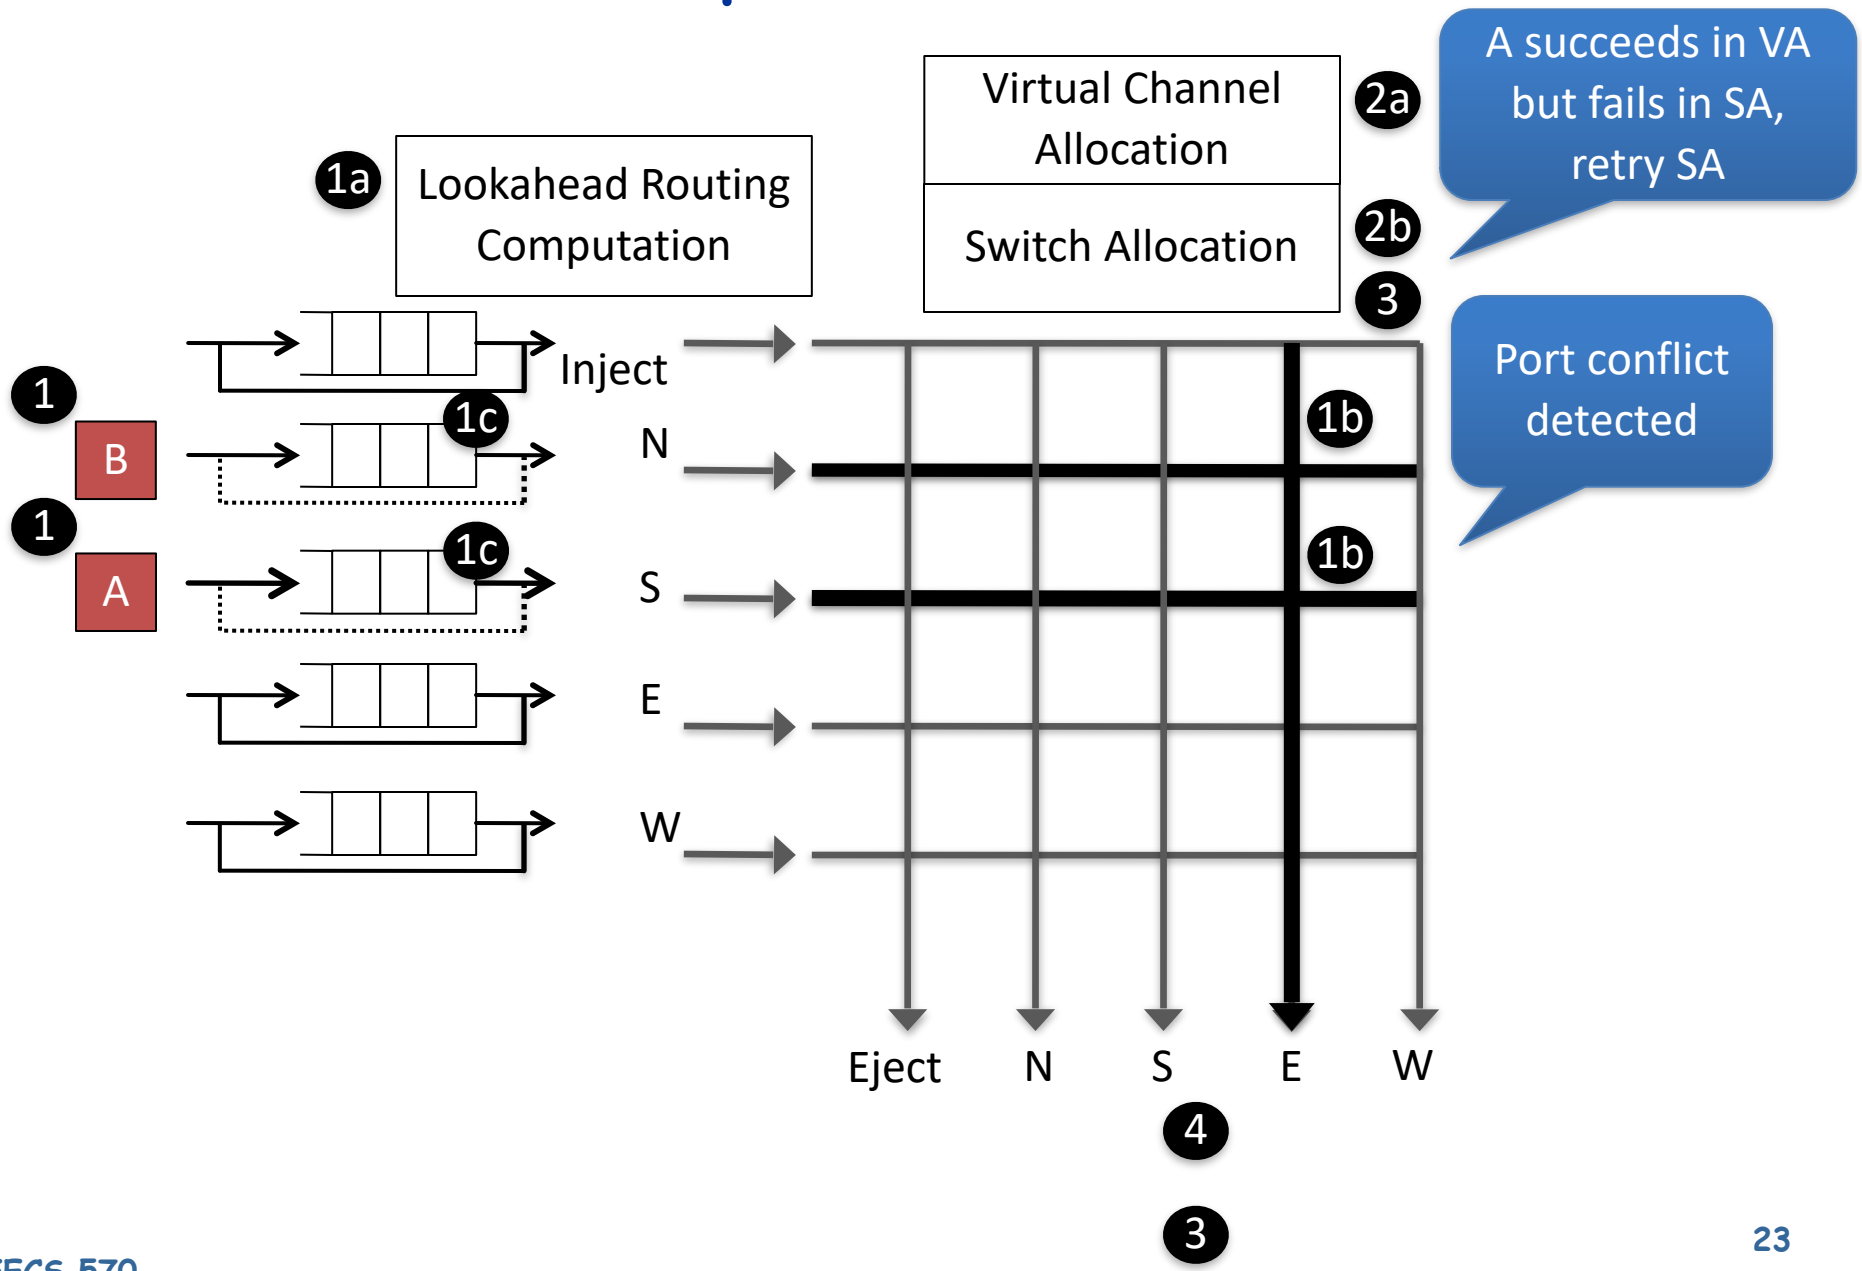

Buffer Organization

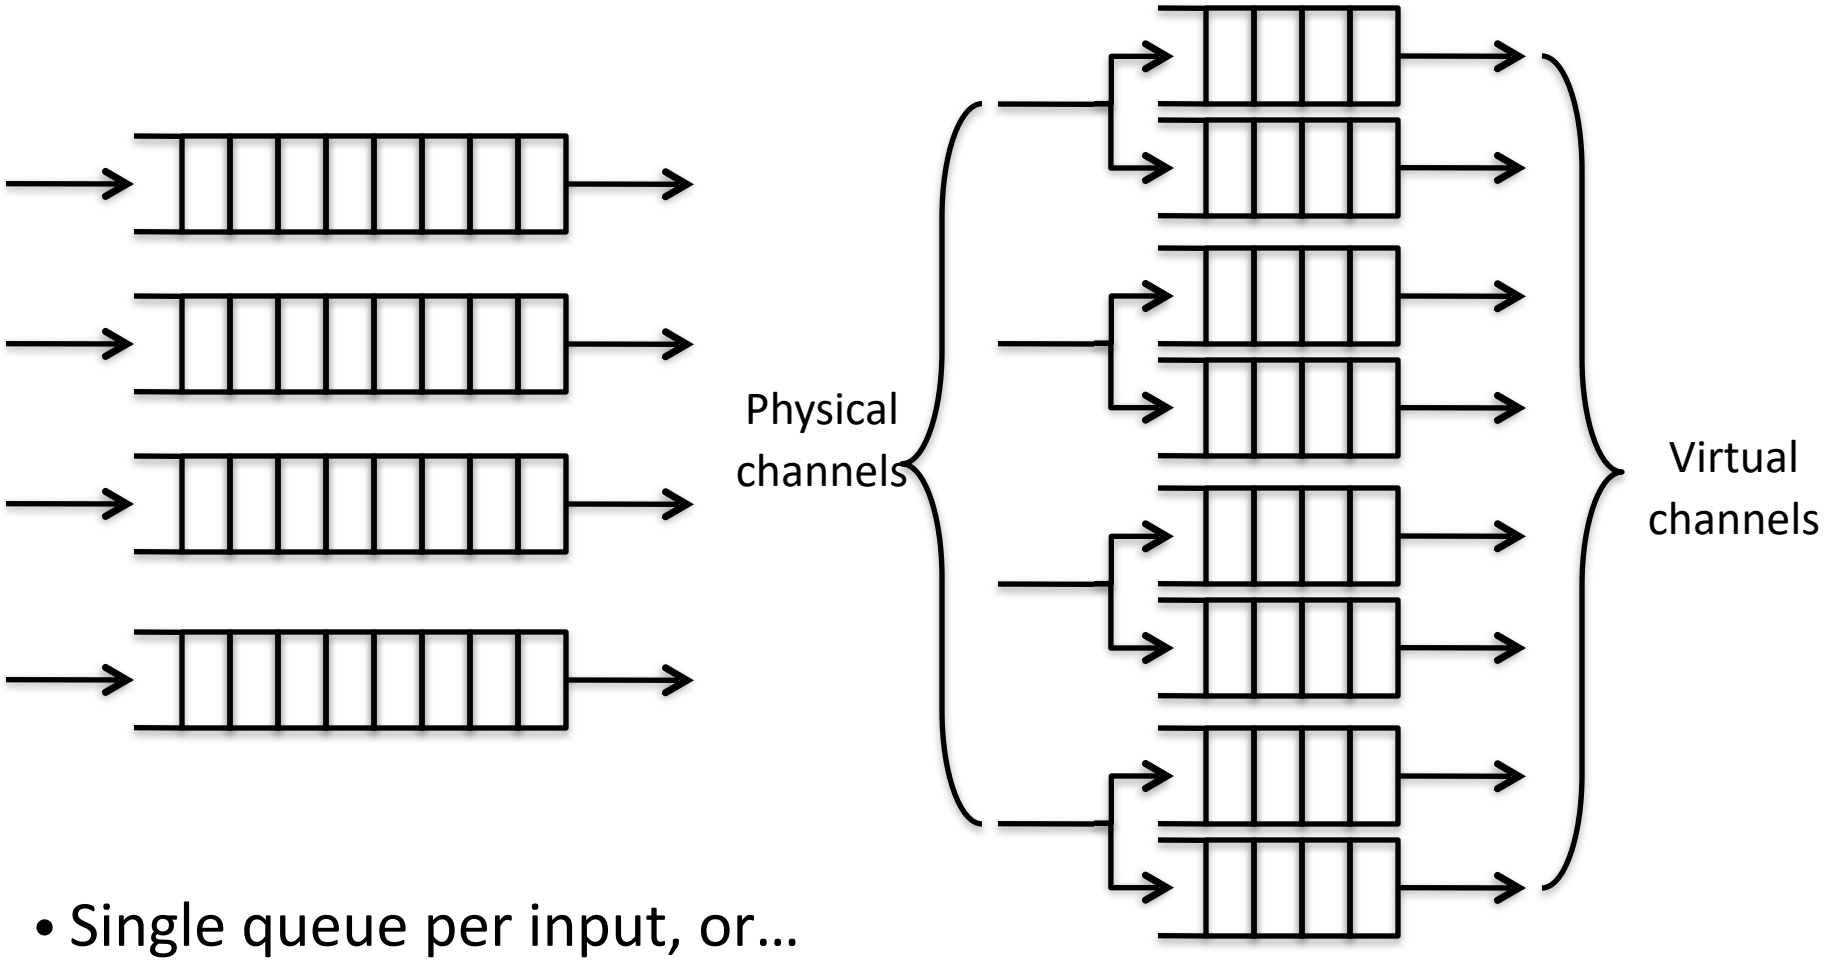

- Single queue per input, or…
- Multiple fixed length queues per physical channel, or…

Buffer Organization

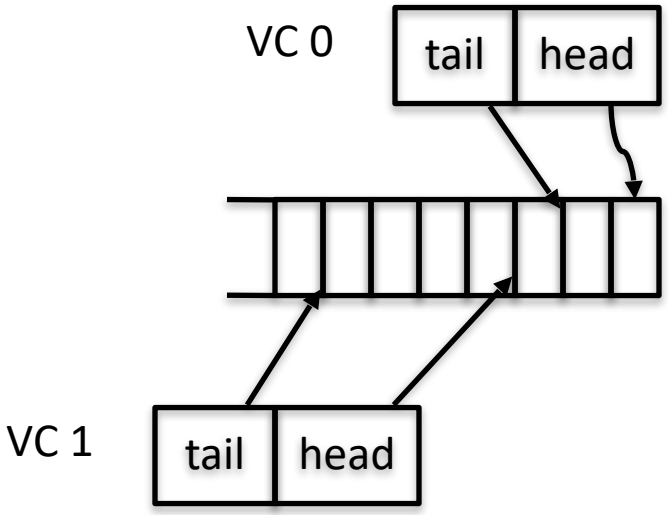

- Multiple variable length queues
  - Multiple VCs share a large buffer
  - Each VC must have minimum 1 flit buffer
    - Prevent deadlock
  - More complex circuitry

Buffer Organization

- Many shallow VCs?
- Few deep VCs?

- More VCs ease HOL blocking
  - More complex VC allocator

- Light traffic
  - Many shallow VCs – underutilized
- Heavy traffic
  - Few deep VCs – less efficient, packets blocked due to lack of VCs

Switch Organization

- Heart of datapath
  - Switches bits from input to output
- High frequency crossbar designs challenging
- Crossbar composed for many multiplexers
  - Common in low-frequency router designs

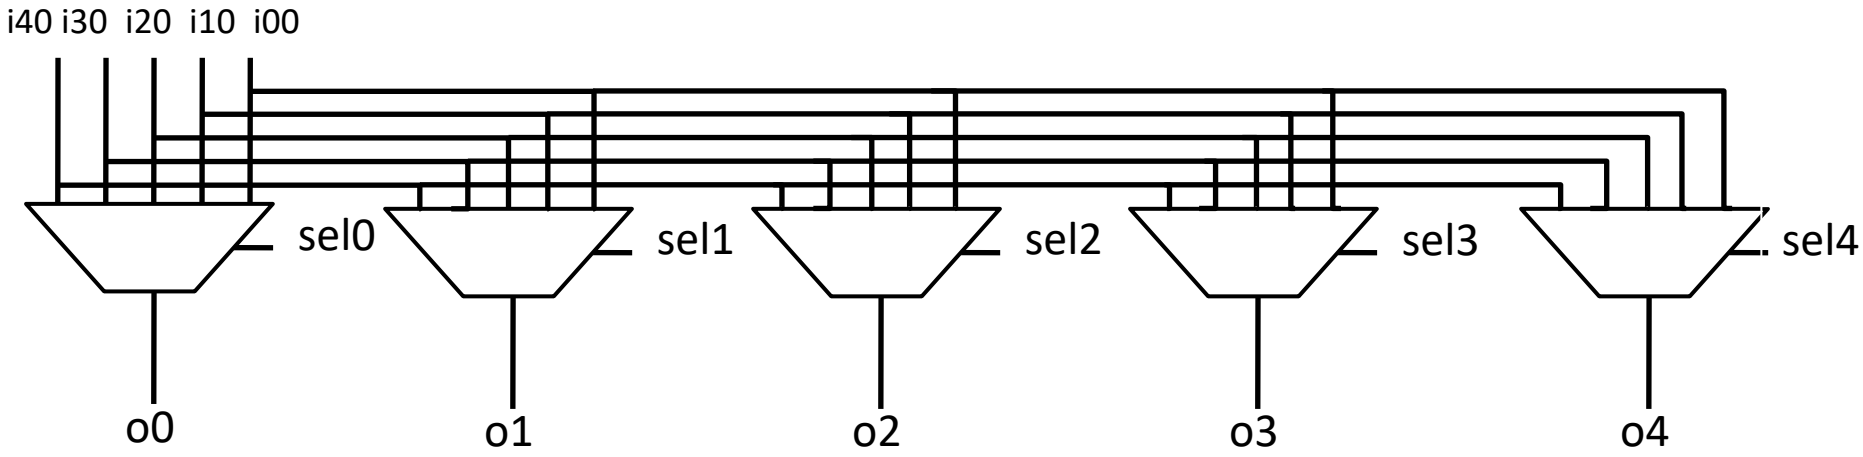

Switch Organization: Crosspoint

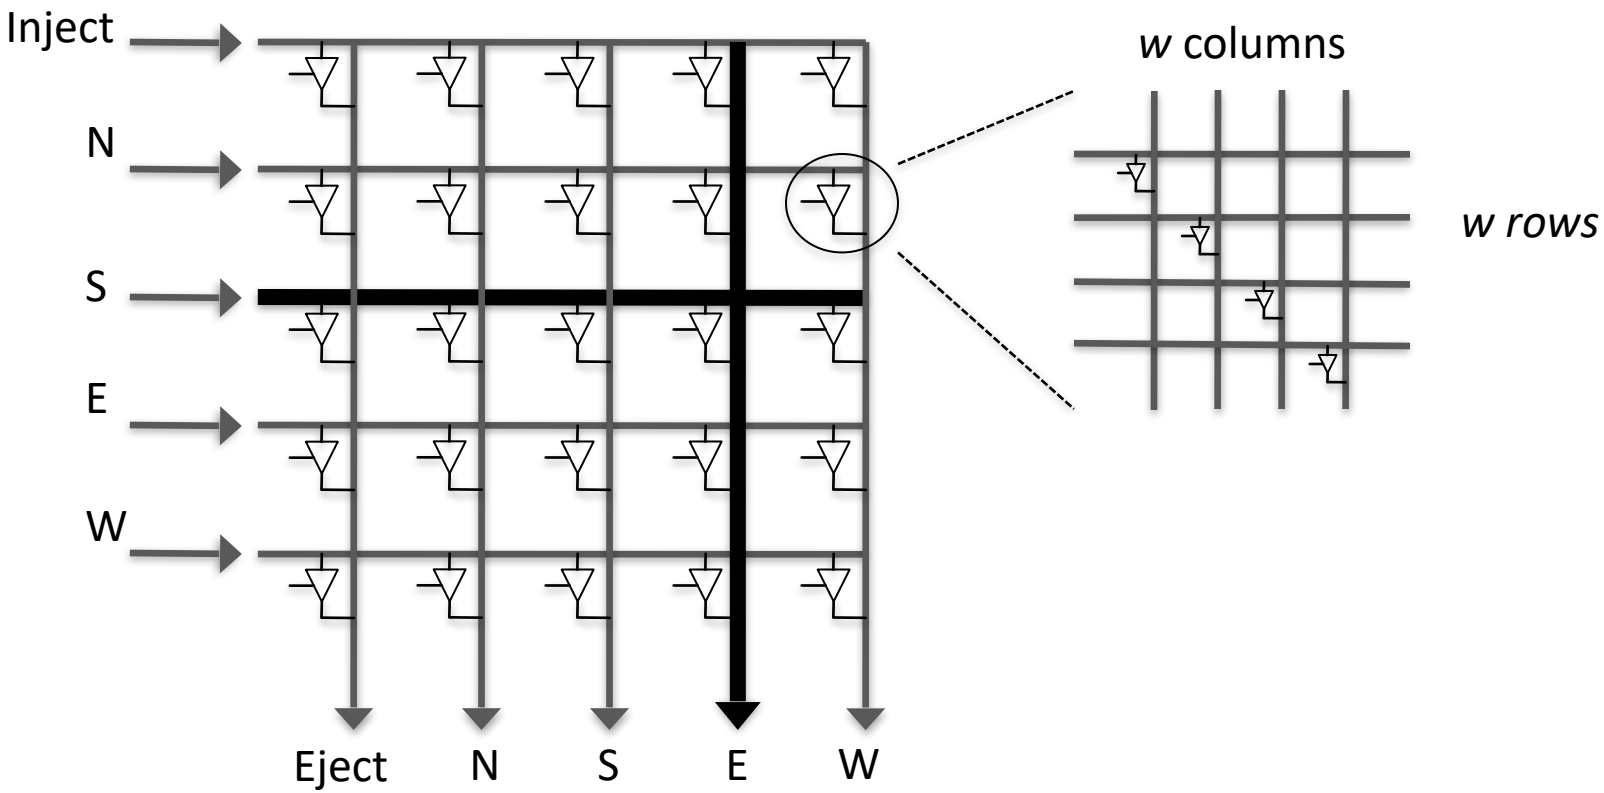

- Area and power scale at $O((pw)^2)$
  - p: number of ports (function of topology)
  - w: port width in bits (determines phit/flit size and impacts packet energy and delay)

Crossbar speedup

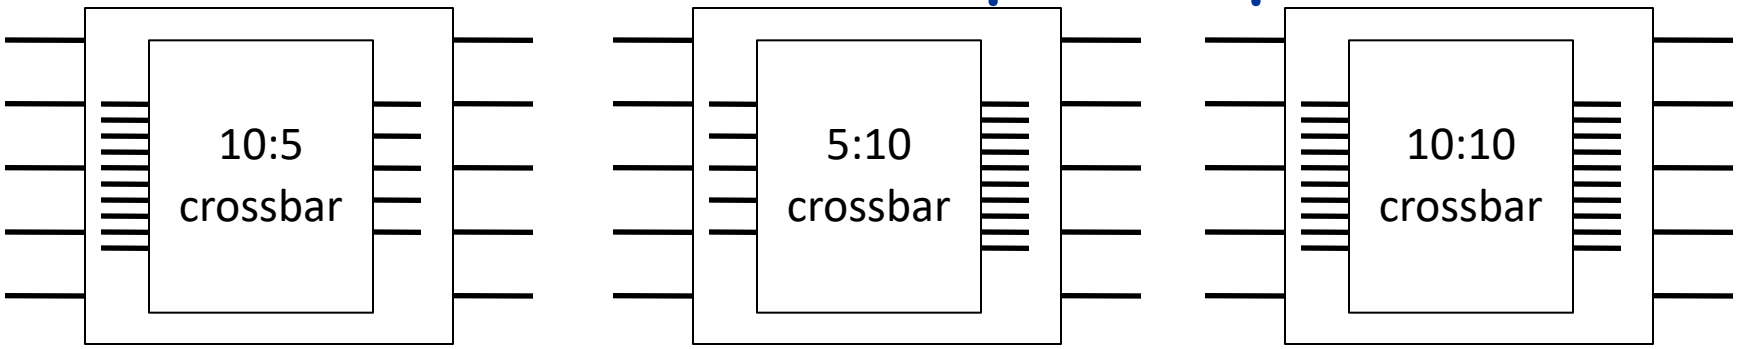

- Refers to # of input and output ports in crossbar relative to the # of router input and output ports
- Increases internal switch bandwidth
- Simplifies allocation or gives better performance with a simple allocator
  - More inputs to select from → higher probability each output port will be matched (used) each cycle
- Output speedup requires output buffers
  - Multiplex onto physical link