


# Lecture 22: Interconnection Networks

---


- Topics: Routing, deadlock, flow control, virtual channels

# Network Topology Examples


---



Grid



Torus



Hypercube

# Routing

---


- Deterministic routing: given the source and destination, there exists a unique route
- Adaptive routing: a switch may alter the route in order to deal with unexpected events (faults, congestion) – more complexity in the router vs. potentially better performance
- Example of deterministic routing: dimension order routing: send packet along first dimension until destination co-ord (in that dimension) is reached, then next dimension, etc.

# Deadlock

---

- Deadlock happens when there is a cycle of resource dependencies – a process holds on to a resource (A) and attempts to acquire another resource (B) – A is not relinquished until B is acquired


# Deadlock Example



Each message is attempting to make a left turn – it must acquire an output port, while still holding on to a series of input and output ports

# Deadlock-Free Proofs


- Number edges and show that all routes will traverse edges in increasing (or decreasing) order – therefore, it will be impossible to have cyclic dependencies
- Example: k-ary 2-d array with dimension routing: first route along x-dimension, then along y




# Breaking Deadlock

---


- Consider the eight possible turns in a 2-d array (note that turns lead to cycles)
- By preventing just two turns, cycles can be eliminated
- Dimension-order routing disallows four turns
- Helps avoid deadlock even in adaptive routing




West-First



North-Last



Negative-First



Can allow  
deadlocks

# Packets/Flits

---

- A message is broken into multiple packets (each packet has header information that allows the receiver to re-construct the original message)
- A packet may itself be broken into flits – flits do not contain additional headers
- Two packets can follow different paths to the destination  
Flits are always ordered and follow the same path
- Such an architecture allows the use of a large packet size (low header overhead) and yet allows fine-grained resource allocation on a per-flit basis



# Flow Control

---

- The routing of a message requires allocation of various resources: the channel (or link), buffers, control state
- Bufferless: flits are dropped if there is contention for a link, NACKs are sent back, and the original sender has to re-transmit the packet
- Circuit switching: a request is first sent to reserve the channels, the request may be held at an intermediate router until the channel is available (hence, not truly bufferless), ACKs are sent back, and subsequent packets/flits are routed with little effort (good for bulk transfers)


# Buffered Flow Control

- A buffer between two channels decouples the resource allocation for each channel

- Packet-buffer flow control: channels and buffers are allocated per packet

- Store-and-forward

- Cut-through



- Wormhole routing: same as cut-through, but buffers in each router are allocated on a per-flit basis, not per-packet

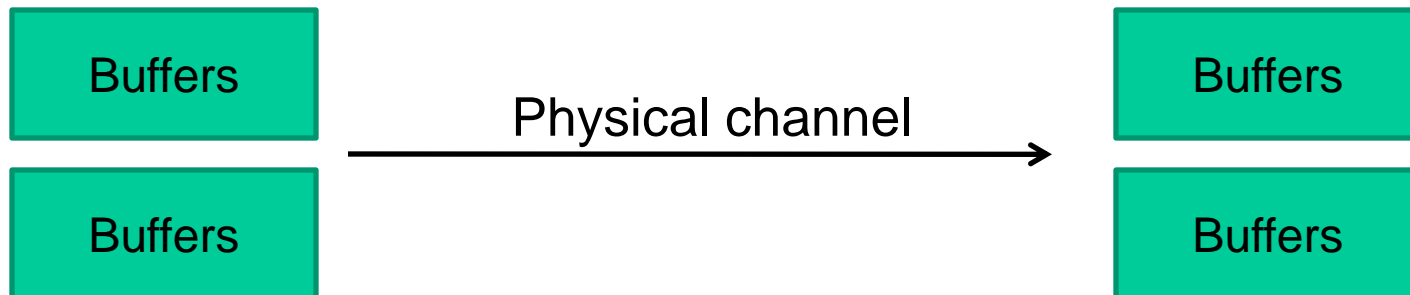
# Virtual Channels

---



Flits do not carry headers. Once a packet starts going over a channel, another packet cannot cut in (else, the receiving buffer will confuse the flits of the two packets). If the packet is stalled, other packets can't use the channel.


With virtual channels, the flit can be received into one of N buffers. This allows N packets to be in transit over a given physical channel. The packet must carry an ID to indicate its virtual channel.



# Example

- **Wormhole:**


A is going from Node-1 to Node-4; B is going from Node-0 to Node-5



(blocked, no free VCs/buffers)

Traffic Analogy:  
B is trying to make a left turn; A is trying to go straight; there is no left-only lane with wormhole, but there is one with VC

- **Virtual channel:**



(blocked, no free VCs/buffers)

# Virtual Channel Flow Control

---

- Incoming flits are placed in buffers
- For this flit to jump to the next router, it must acquire three resources:
  - A free virtual channel on its intended hop
    - We know that a virtual channel is free when the tail flit goes through
  - Free buffer entries for that virtual channel
    - This is determined with credit or on/off management
  - A free cycle on the physical channel
    - Competition among the packets that share a physical channel


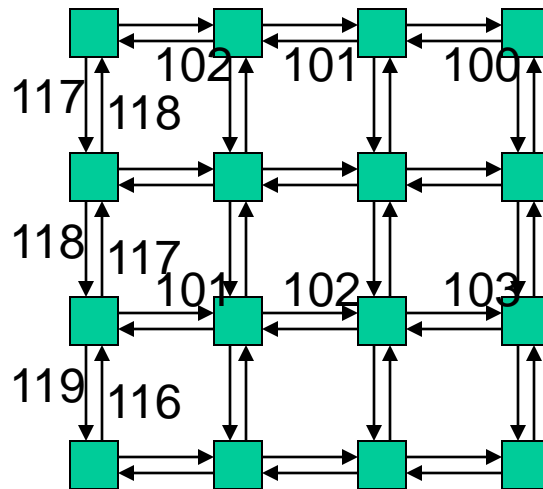

# Buffer Management

---

- Credit-based: keep track of the number of free buffers in the downstream node; the downstream node sends back signals to increment the count when a buffer is freed; need enough buffers to hide the round-trip latency
- On/Off: the upstream node sends back a signal when its buffers are close to being full – reduces upstream signaling and counters, but can waste buffer space

# Deadlock Avoidance with VCs

- VCs provide another way to number the links such that a route always uses ascending link numbers



- Alternatively, use West-first routing on the 1<sup>st</sup> plane and cross over to the 2<sup>nd</sup> plane in case you need to go West again (the 2<sup>nd</sup> plane uses North-last, for example)

# Title

---

- Bullet