NoC

Po-An Tsai

6.823S15 Recitation

NoC

- Network-on-chip is about how elements (cores, cache banks, memory controller, I/O controller etc.) communicate with each other

- Important when you have a large system

Knights Landing Processor Architecture

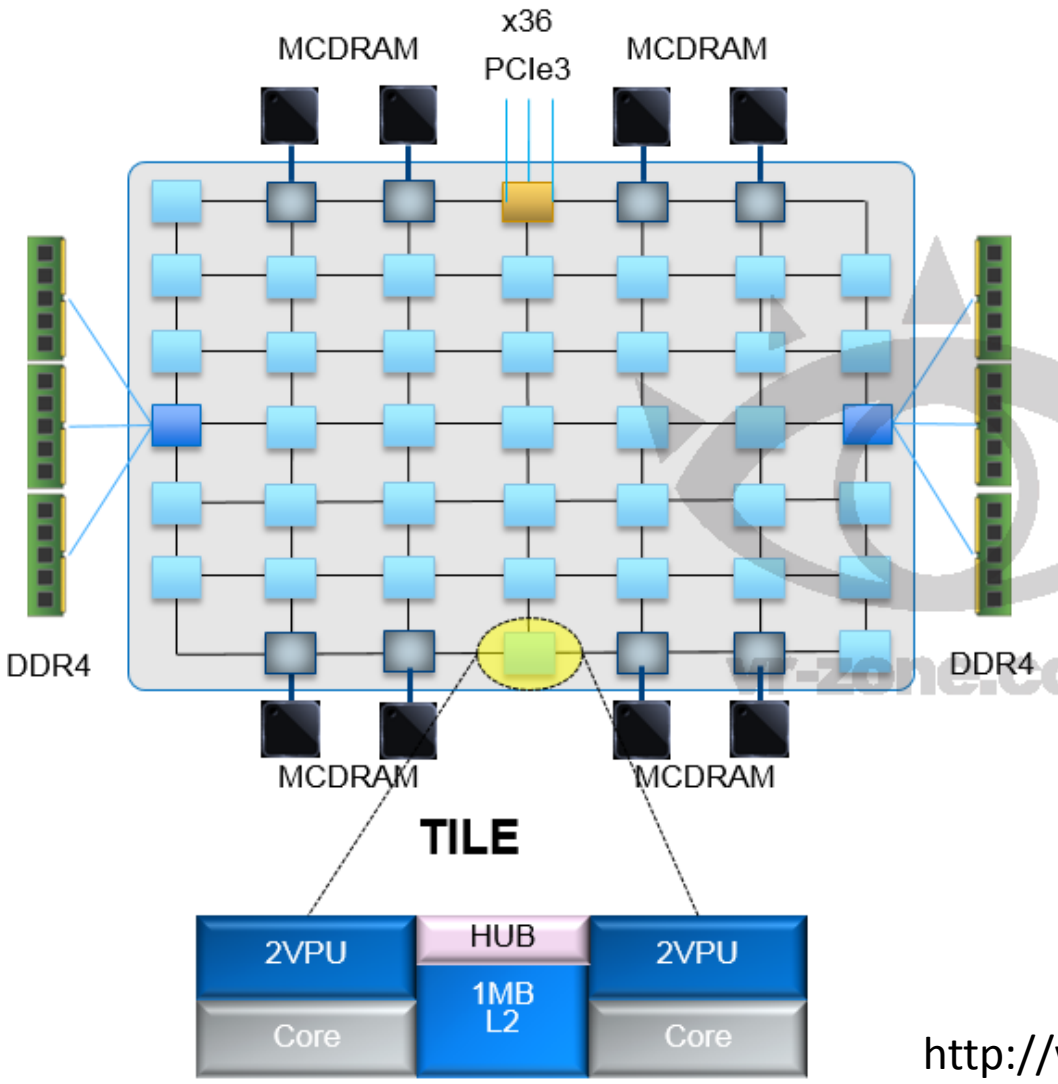

- Up to 72 Intel Architecture cores based on Silvermont (Intel® Atom processor)
  - Four threads/core
  - Two 512b vector units/core
  - Up to 3x single thread performance improvement over KNC generation
- Full Intel® Xeon processor ISA compatibility through AVX-512 (except TSX)
- 6 channels of DDR4 2400 MHz -up to 384GB
- 36 lanes PCI Express* Gen 3
- 8/16GB of high-bandwidth on-package MCDRAM memory >500GB/sec
- 200W TDP

http://www.extremetech.com/wp-content/uploads/2013/11/KNL13.png

Topology

- How different nodes connect to each other
  - Ring
  - Mesh/Torus
  - Tree
- Important properties
  - Diameter
  - Avg. distance
  - Bisection bandwidth
  - Links (overhead)

Topology

Problem 8.2

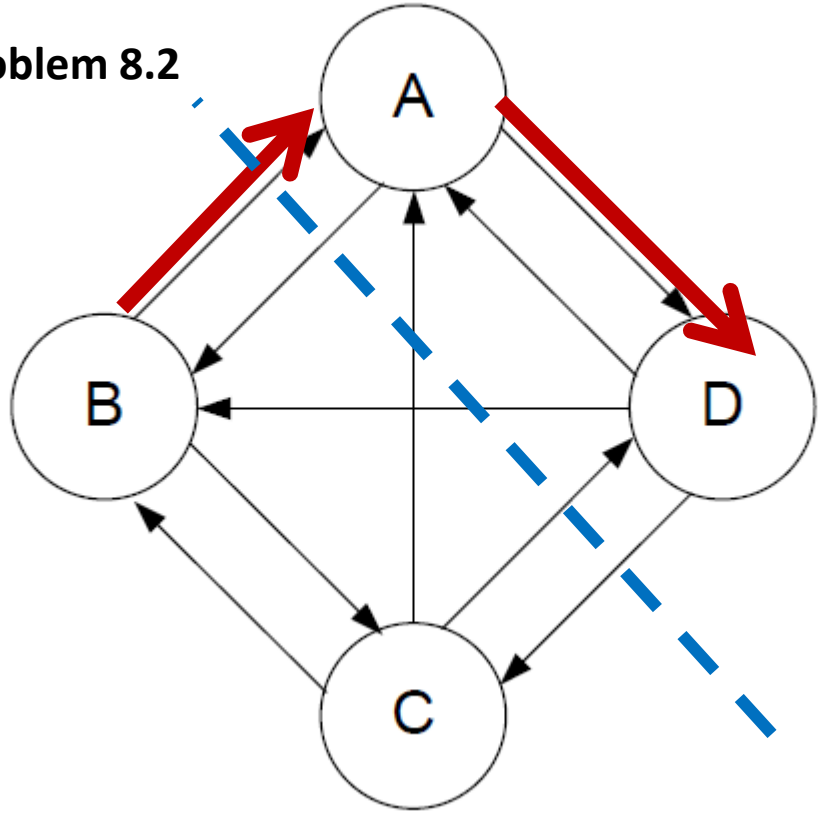

Diameter? 2

Average Hop Count?

(AB+AC+AD+BA+BC+BD+CA+CB+CD+DA+DB+DC)/12 = 7/6

Bisection Bandwidth?

4 if all links bi-directional

Flow control

- How messages are forwarded from src to dst

- Buffered/ Bufferless
  - Wormhole is the most common one
  - but there is head-of-line blocking problem

Head-of-line Blocking in Street Network

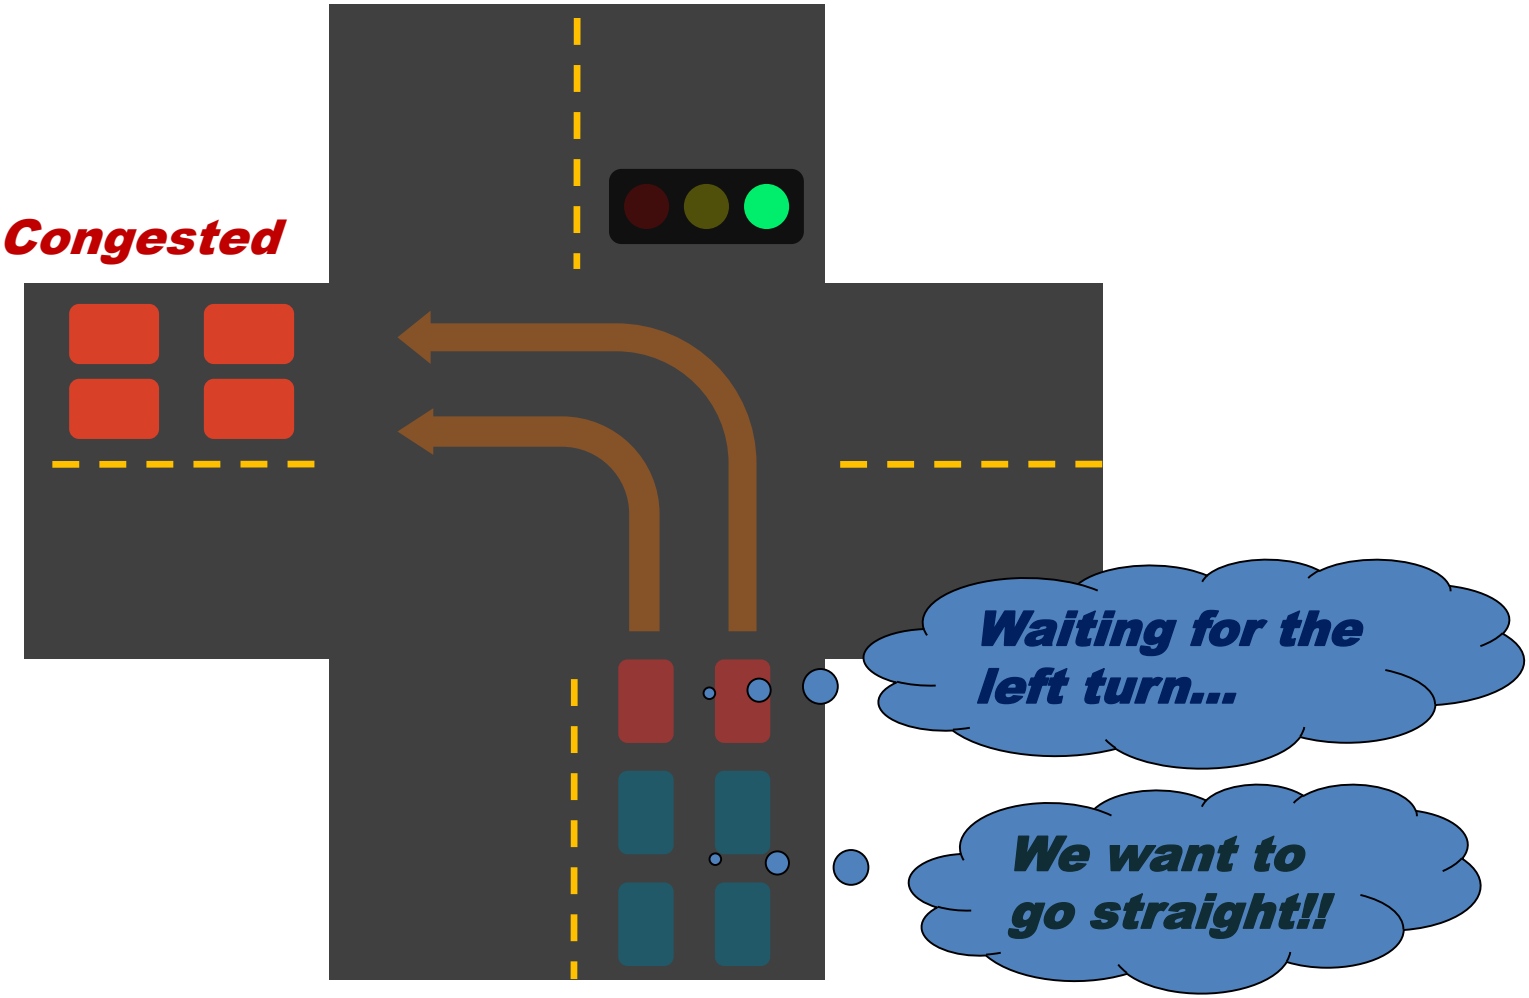

This is why we have such lanes as “straight only” or “left turn only”

Congestion & HoL Blocking

- Head-of-Line (HoL) Blocking

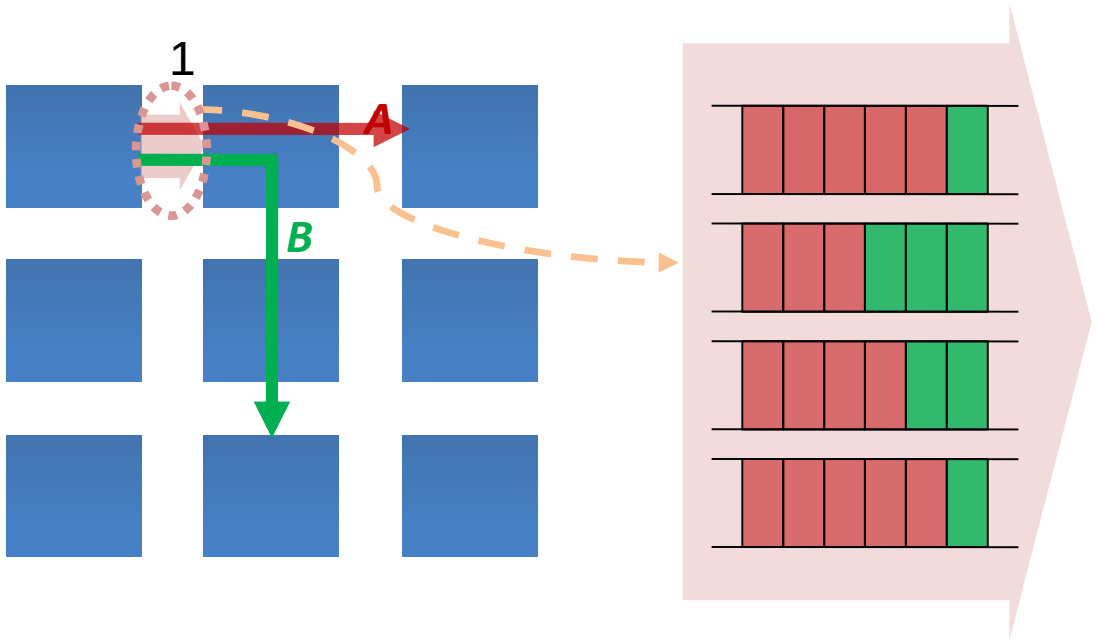

Congestion & HoL Blocking

- Head-of-Line (HoL) Blocking

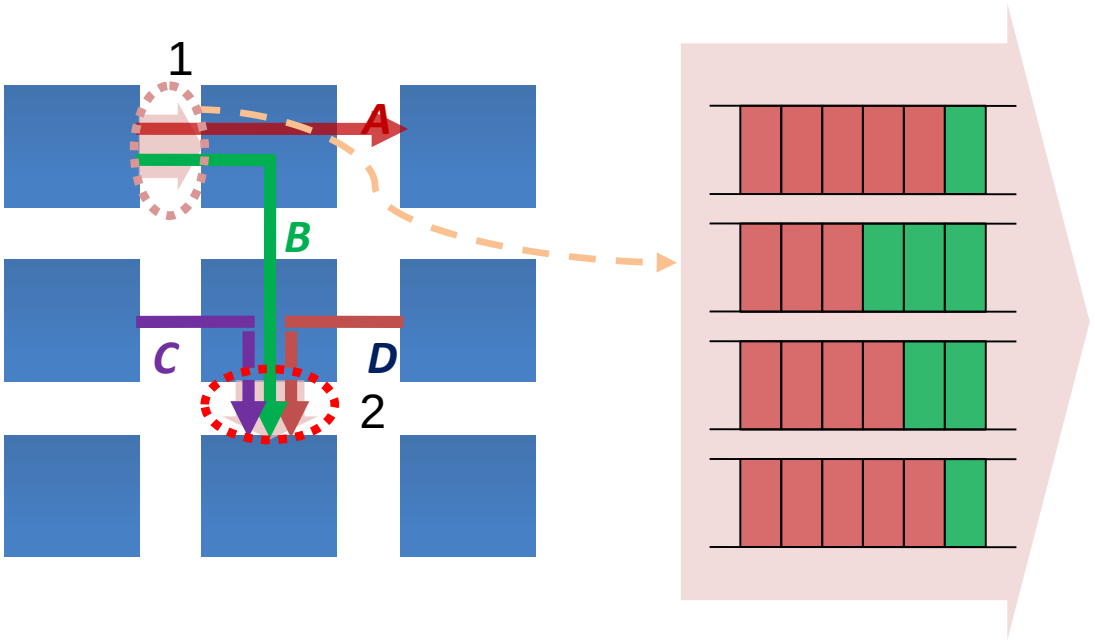

Solution: Virtual Channels

Congestion & HoL Blocking

- Head-of-Line (HoL) Blocking

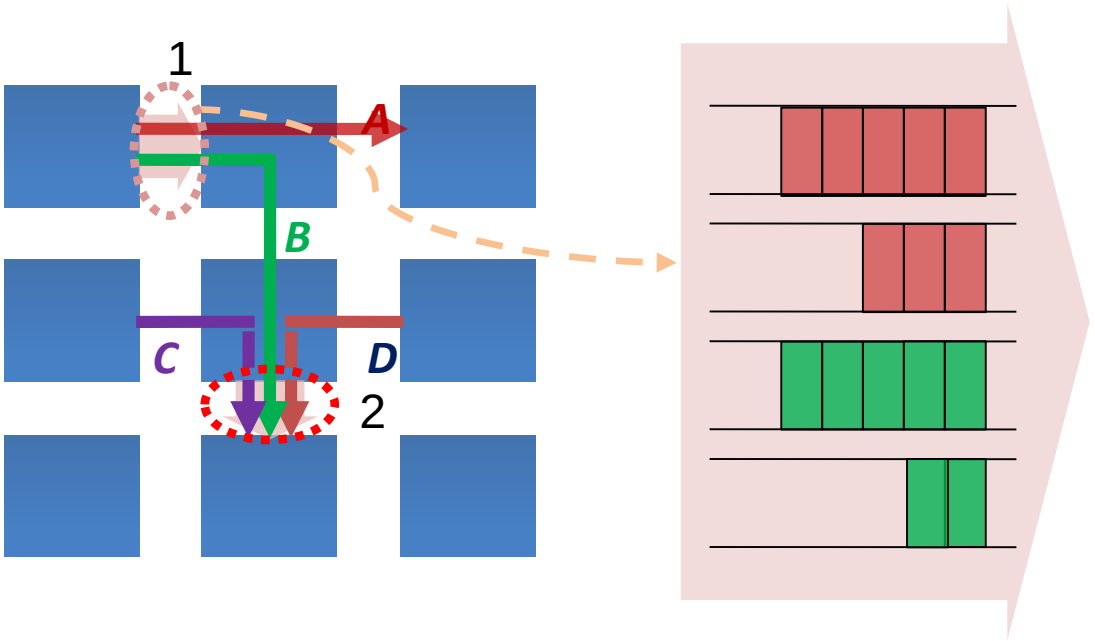

Solution: Virtual Channels

Routing

- What is the path between src and dst
  - Use mesh as example here
- Choose a path so that the message can arrive faster

- Choose a path to ensure no deadlock/livelock

Routing: Deadlock

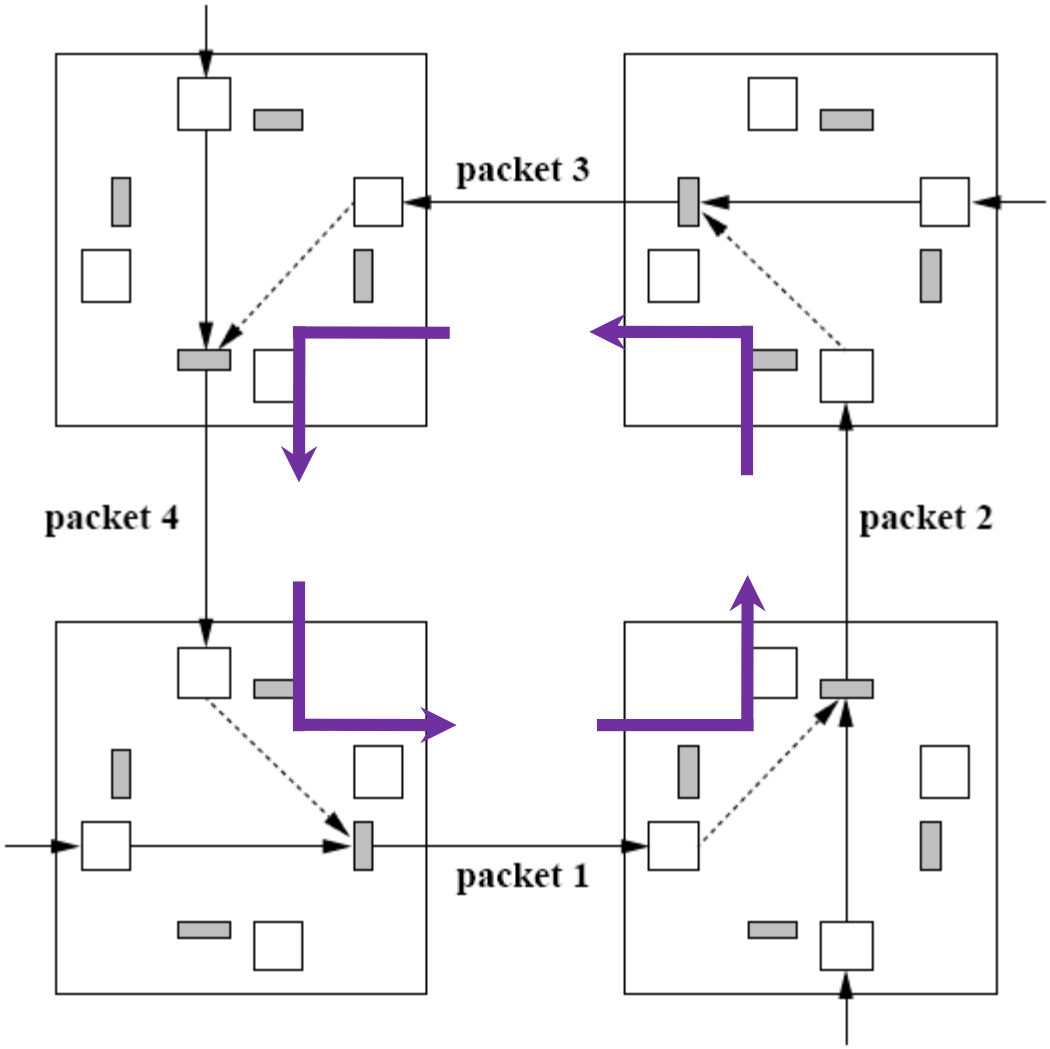

Turn Model

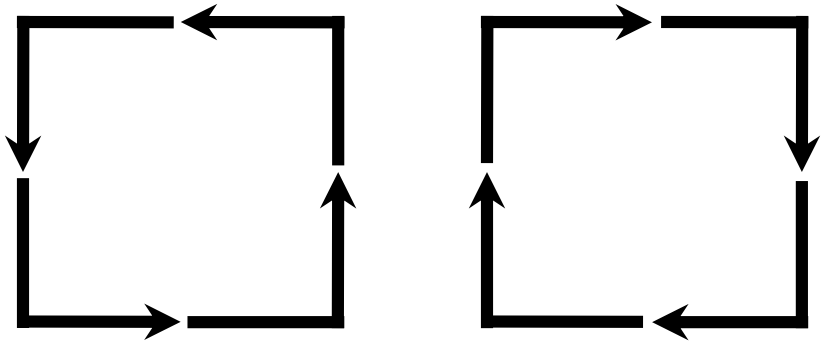

The eight possible turns and cycles in a 2D mesh

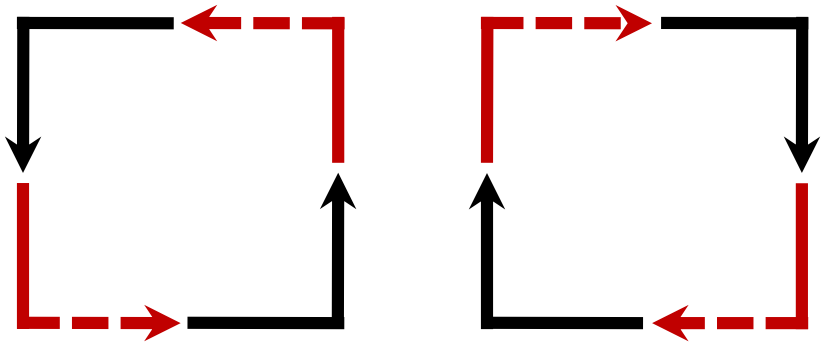

Only four turns are allowed in the XY routing algorithm

Turn Model

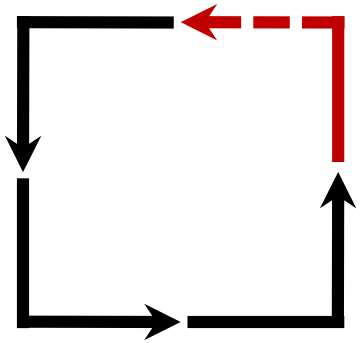

AND

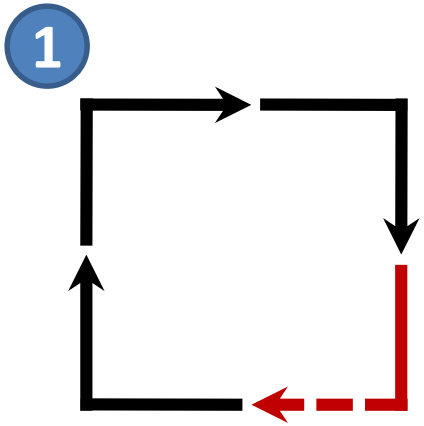

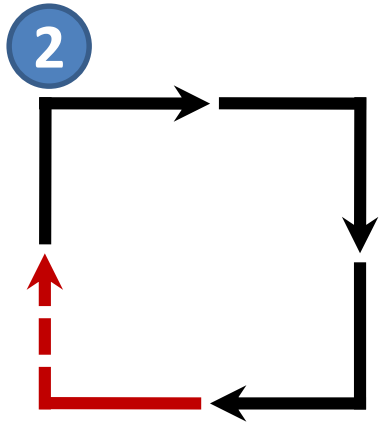

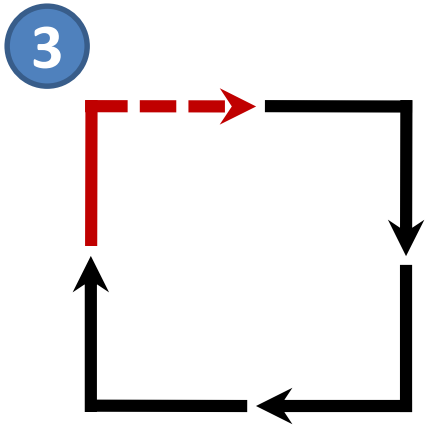

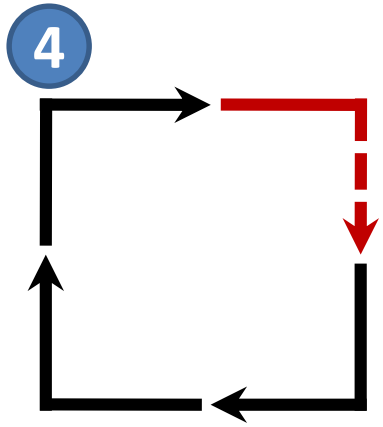

Turn Model

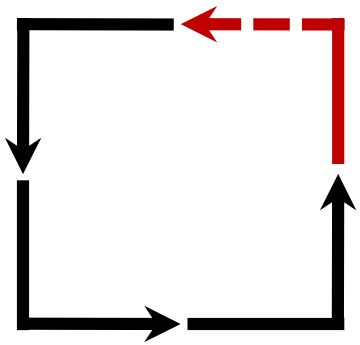

AND

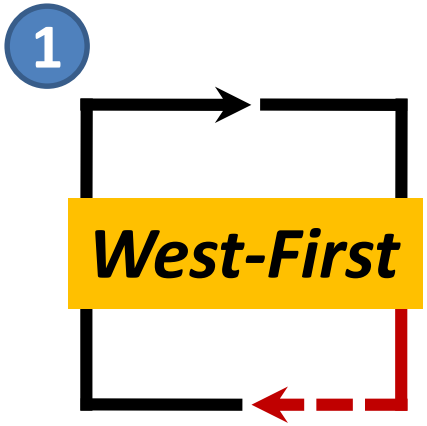

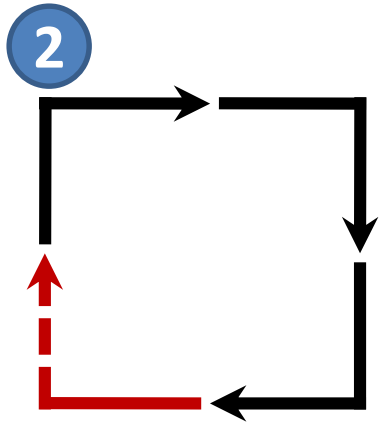

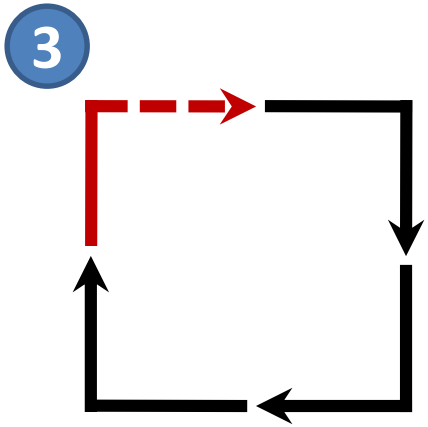

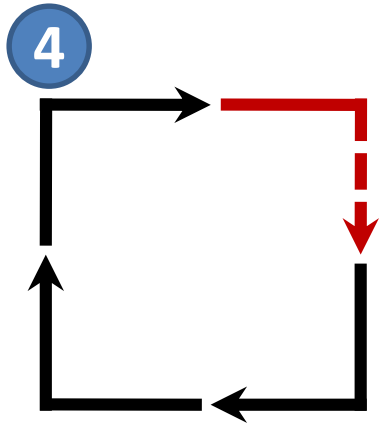

Turn Model

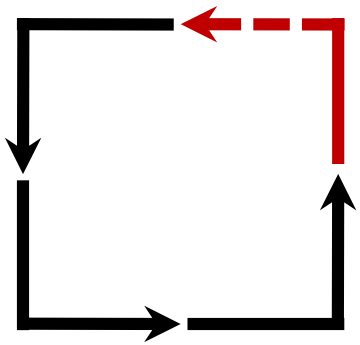

AND

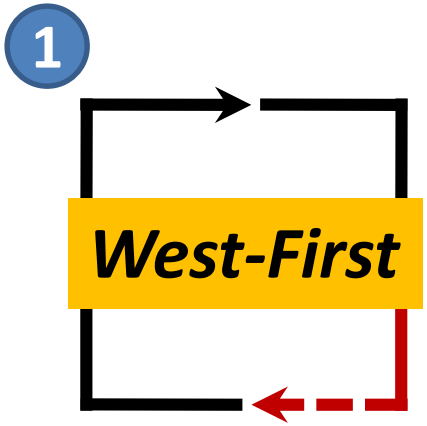

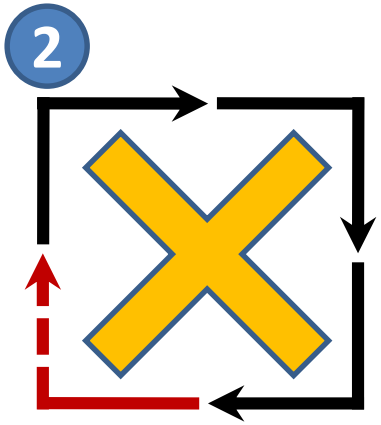

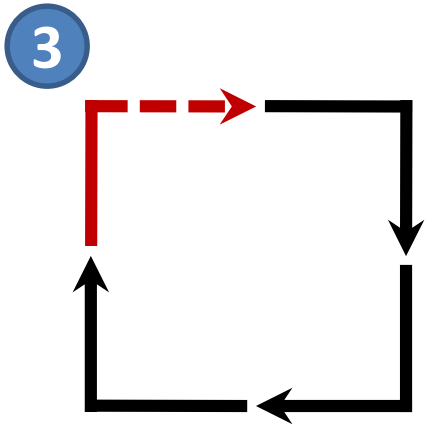

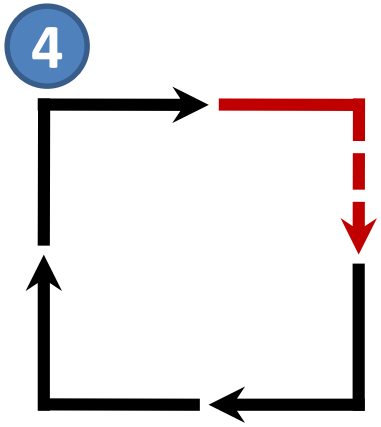

Turn Model

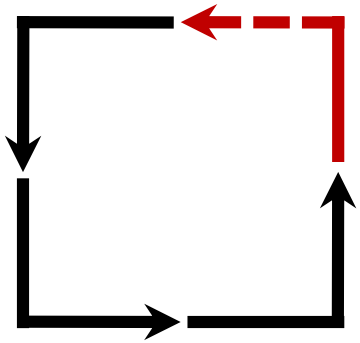

AND

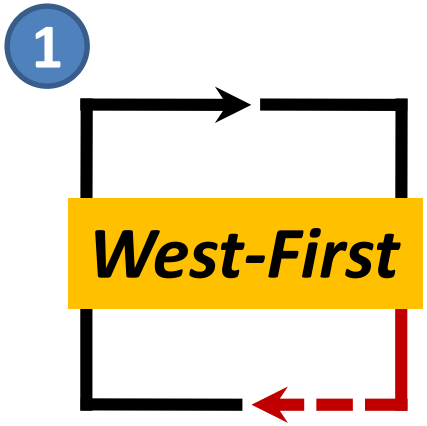

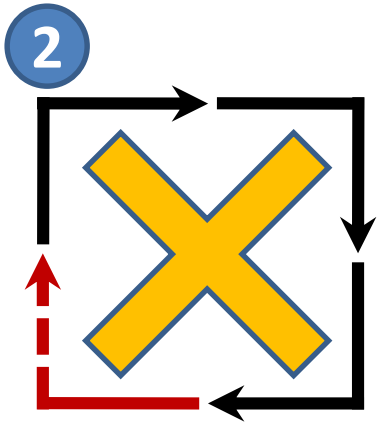

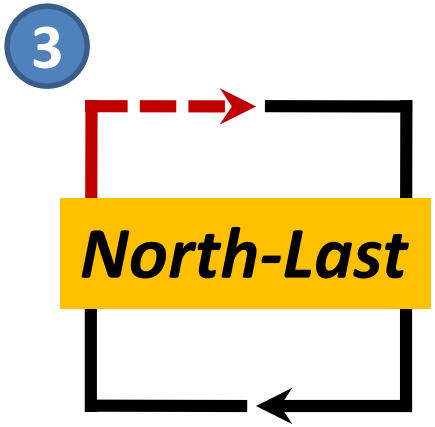

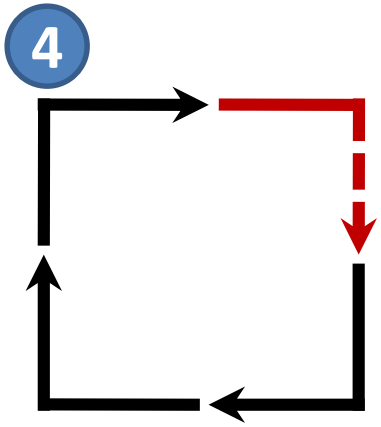

Turn Model

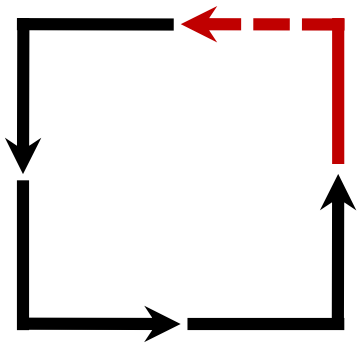

AND

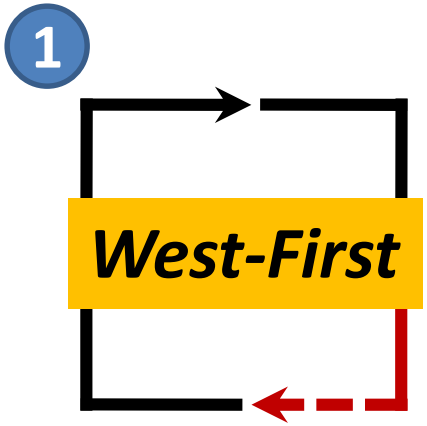

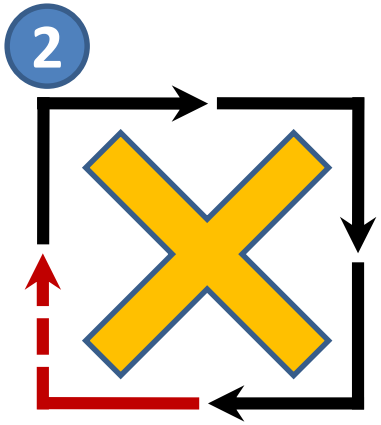

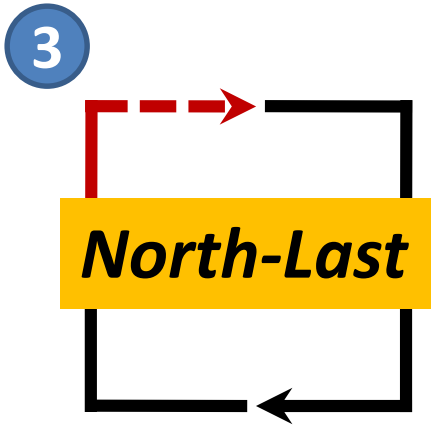

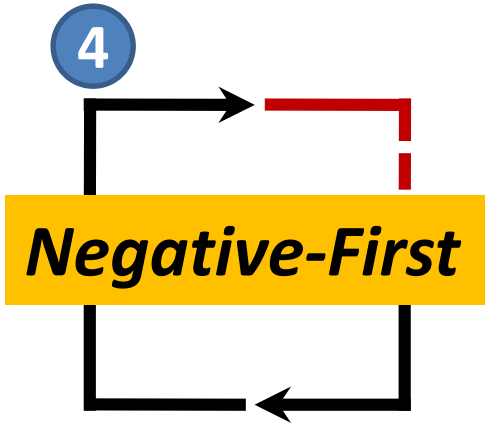

Channel dependency graph (CDG)

Problem 8.2

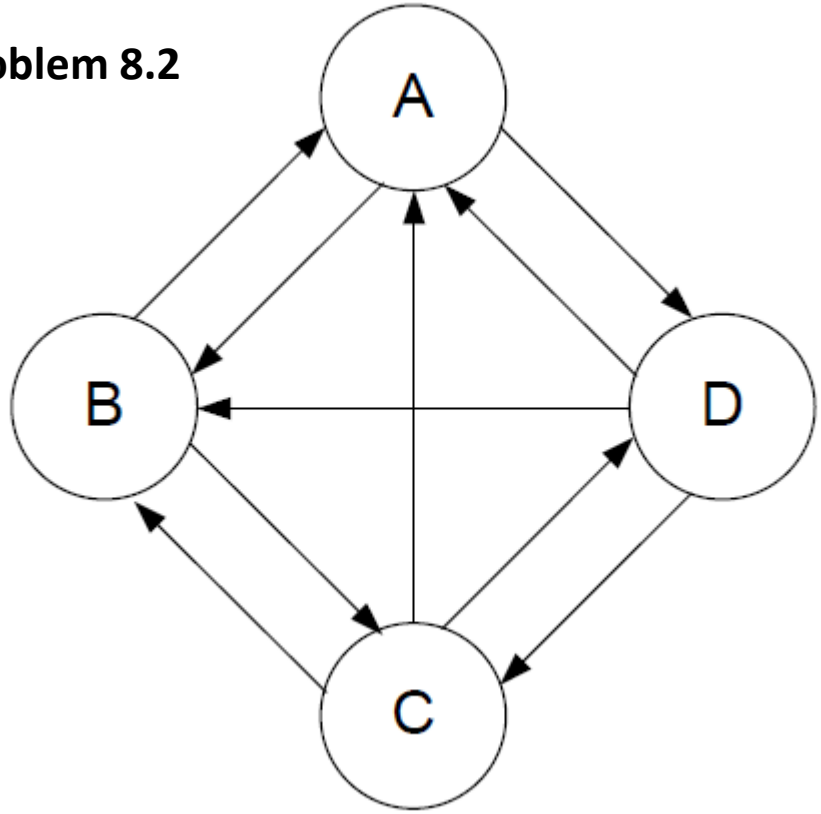

8.2.B: CDG

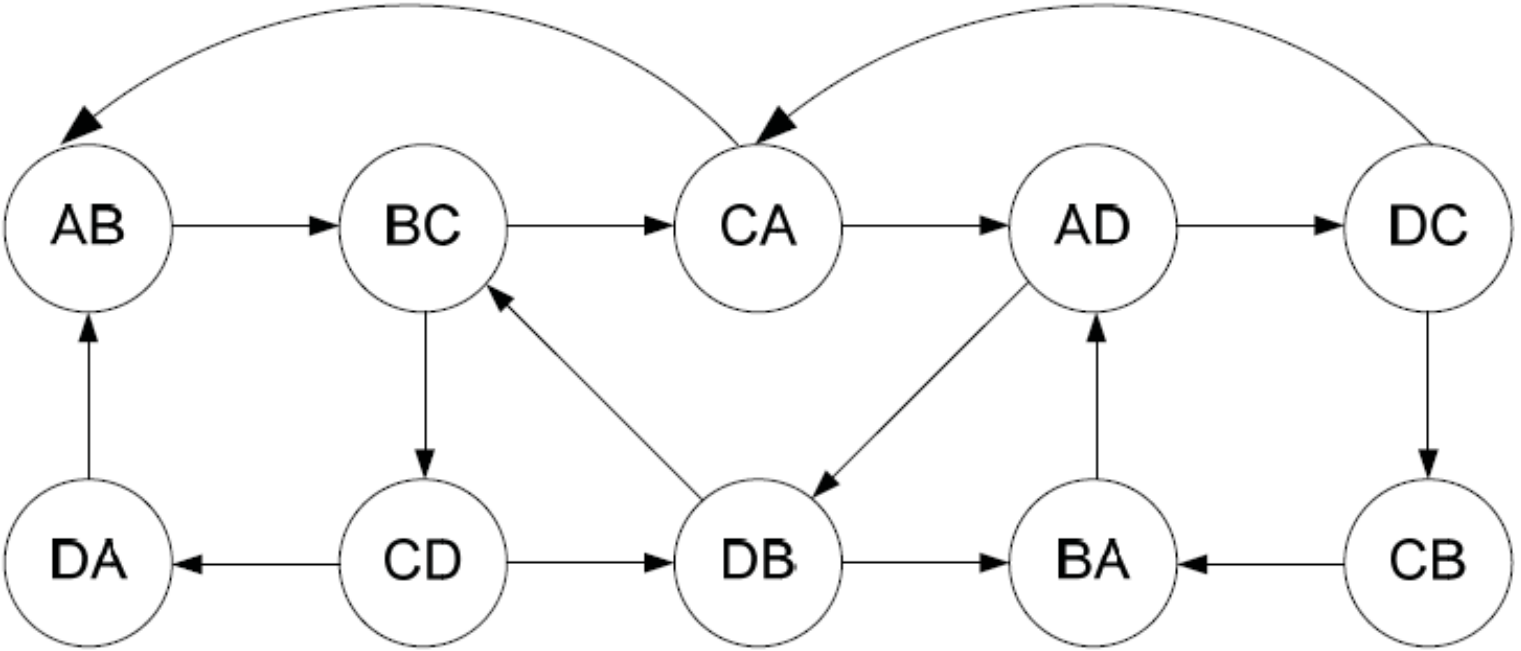

Deadlock free? No

8.2.C: Minimal Routing

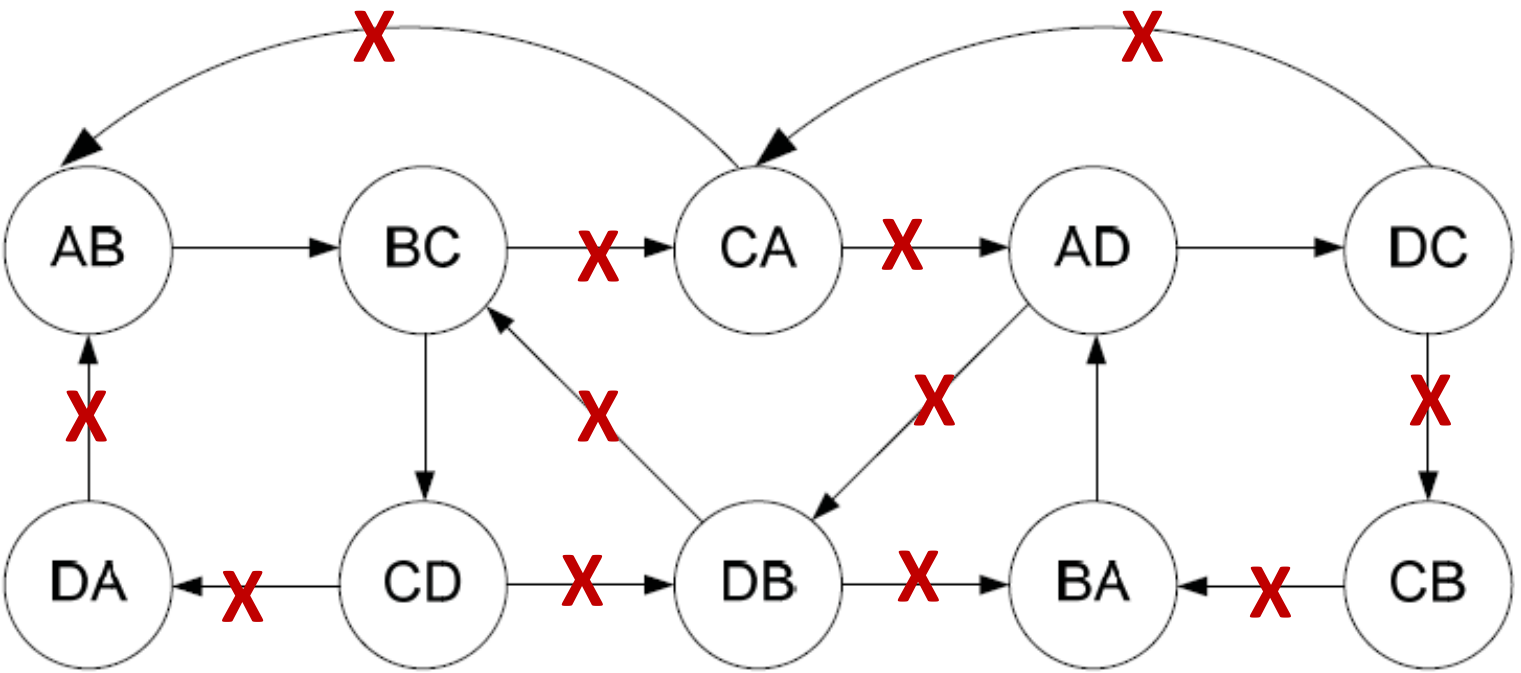

Deadlock free? Yes

Dimension-Order Routing (DOR)

- Approaches in one dimension first, then in the other

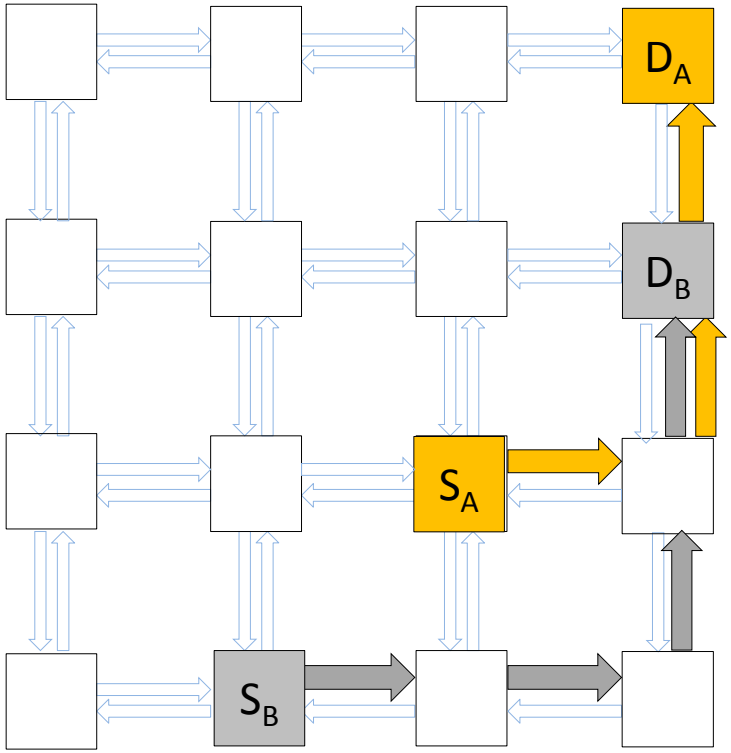

DOR (XY)

- Bandwidth
  No path diversity
- Latency
  Minimal routing
- Deadlock Prevention
  Deadlock-free with 1 VC

Valiant

- Uses one random intermediate node per each packet

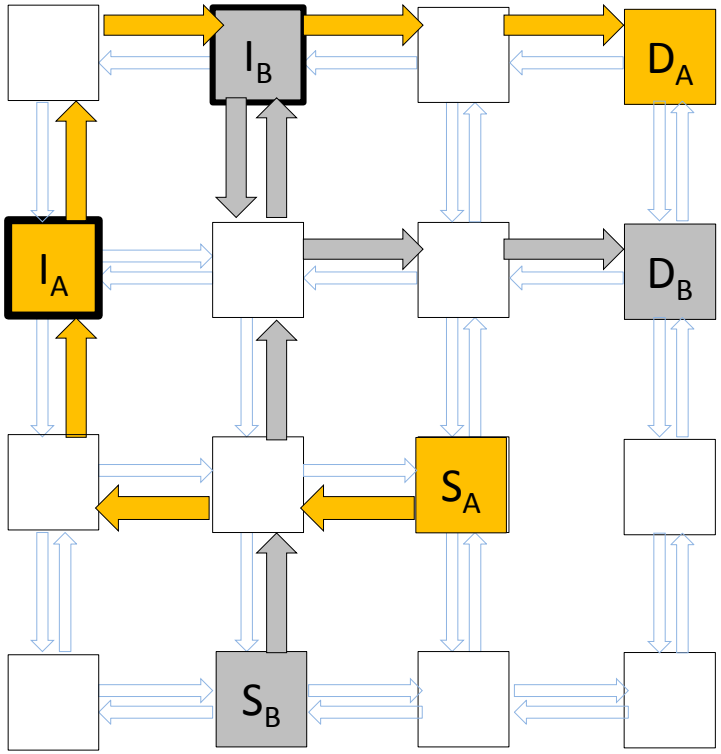

Valiant (XY/YX)

- Bandwidth
  Wide path diversity
- Latency
  Poor latency
- Deadlock Prevention
  Deadlock-free with >= 2 VCs
  - each phase should use different VCs

n-phase ROMM

- ***n-1*** random intermediate node(s) only in MBR (Minimum Bounding Rectangle)

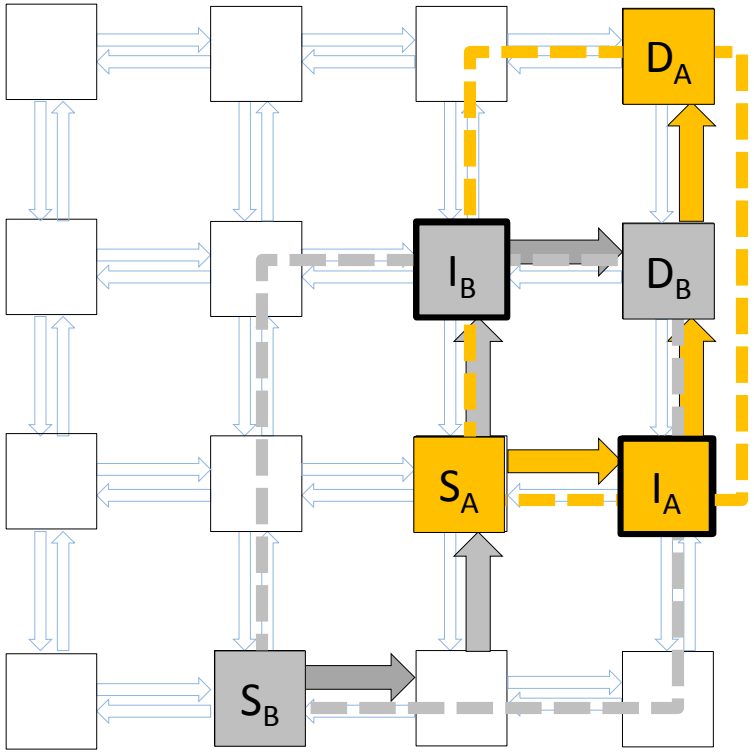

2-phase ROMM (XY/YX)

- Bandwidth
  More path diversity than DOR
  Limited by the value of n
- Latency
  Minimal routing
- Deadlock Prevention
  Deadlock-free with >= n VCs
  - each phase should use different VCs

Routing and Performance

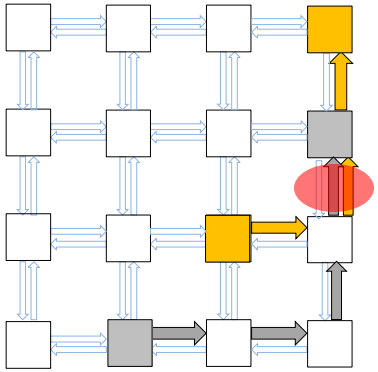

DOR

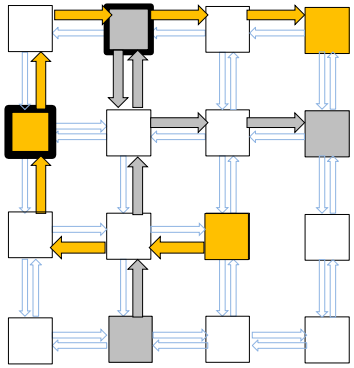

Valiant

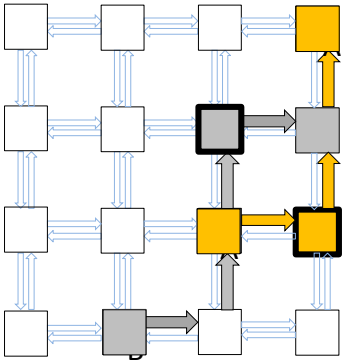

ROMM

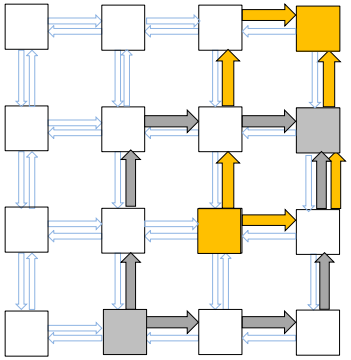

O1TURN

- Depend on traffic patterns
- **In general, path diversity helps lower congestions due to load balancing.**

The end

Next time:

Router architecture

Cache coherence