

ECE 4750 Computer Architecture, Fall 2016

T06 Fundamental Network Concepts

School of Electrical and Computer Engineering
Cornell University

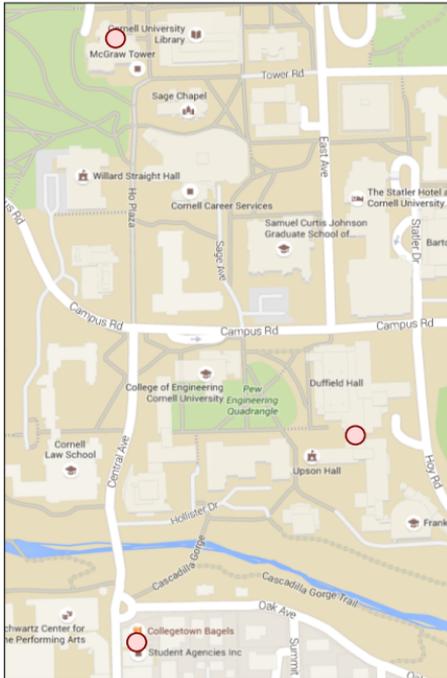
revision: 2016-10-16-22-06

1	Network/Roadway Analogy	3
1.1.	Running Errands	4
1.2.	Network Technology	5
1.3.	Networks in Computer Architecture	6
2	Network Topology	8
2.1.	Single-Stage Bus Topology	8
2.2.	Single-Stage Crossbar Topology	9
2.3.	Multi-Stage Butterfly Topology	10
2.4.	Multi-Stage Torus Topology	12
2.5.	Terminology	14
3	Network Routing	18
3.1.	Oblivious Deterministic Routing	18
3.2.	Oblivious Non-Deterministic Routing	18
3.3.	Adaptive Routing	19
3.4.	Deadlock	20

4 Analyzing Network Performance	21
4.1. Traffic Patterns	21
4.2. Ideal Throughput	23
4.3. Zero-Load Latency	25
4.4. Comparing Topologies	28
4.5. Comparing Routing Algorithms	29

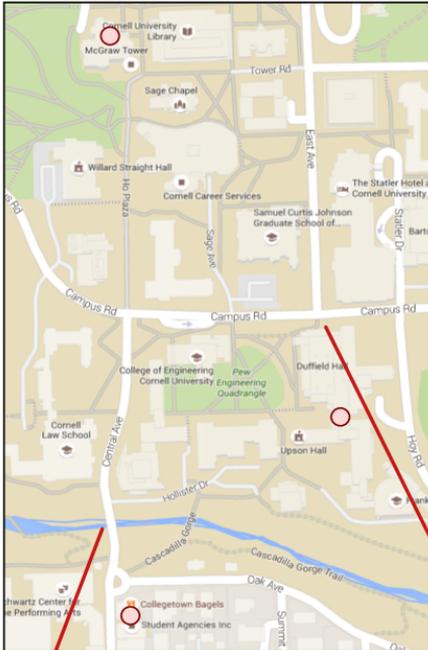
1. Network/Roadway Analogy

Our goal is to run some errands around town using our bike. Assume we are studying in the Duffield Atrium and we need to: (1) do some laundry in Collegetown (maybe pickup some coffee?), and (2) pick up some books about computer architecture from the library.

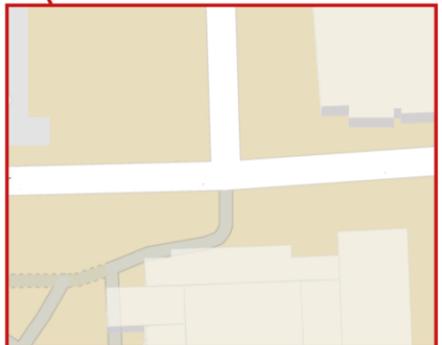


- Duffield Atrium
- Coffee/Laundry
- Library

1.1. Running Errands



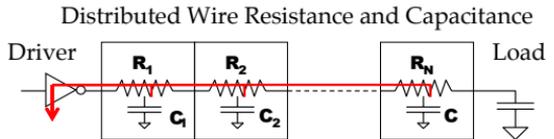
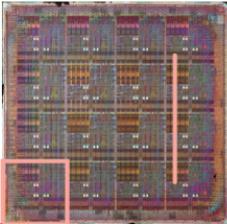
- Network Topology
 - Arrangement of roads and intersections to interconnect sources and destinations
 - Wide vs. narrow roads
 - Long vs. short roads
 - Small vs. large intersections
- Network Routing
 - Path from source to destination along roads and intersections
 - Short vs. long paths
 - Common vs. rare paths
- Network Microarchitecture
 - Managing long line of bikes and cars on road
 - Managing many cars and bikes at same intersection



1.2. Network Technology

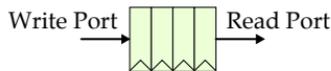
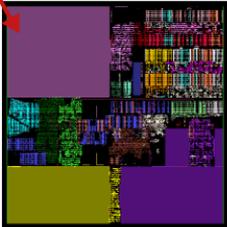
	AWG24 Twisted Pair	PCB Trace	On-Chip M6 Wire in 0.18 μ m
Resistance	0.08 Ω /m	5 Ω /m	40 k Ω /m
Inductance	400 nH/m	300 nH/m	4 μ H/m
Capacitance	40 pF/m	100 pF/m	300 pF/m
Data Rate	\approx Gb/s	\approx Gb/s	\approx Gb/s
Critical Length	1m	10cm	<1cm
Pitch	\approx mm	<mm	< μ m

On-Chip Wires



Because both wire resistance and wire capacitance increase with length, wire delay grows quadratically with length

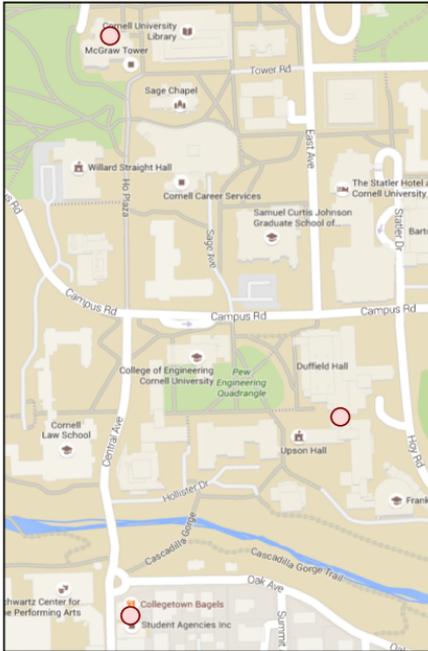
On-Chip Buffers



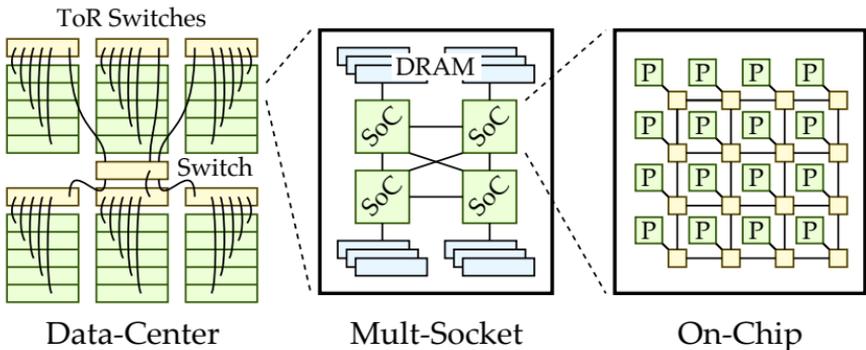
On-chip buffers are 1r1 FIFOs implemented using either a register file or SRAM

On-chip network technology constraints very different from off-chip technology constraints

1.3. Networks in Computer Architecture

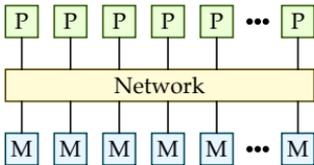


- Network Topology
 - Arrangement of channels and routers to interconnect sources and destinations
 - Roads = channels (implemented with cables, traces, wires)
 - Width of road = channel bw
 - Intersections = routers
 - Size intersection = router radix
- Network Routing
 - Path from source to destination along channels and routers
 - Short vs. long paths = minimal vs. non-minimal paths
 - Common vs. rare paths = channel congestion
- Topology is constrained by packaging (geography)

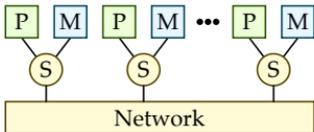


Network Transactions

We will use processor-memory networks as a driving example throughout the course. Processor-memory networks allow many processors to perform memory read/write transactions on many memories.



Dance-Hall Organization



Integrated-Node Organization

- Message = Logical unit of data transfer provided by network interface
- Packet = Unit of routing within a network
- Flit = Smallest unit of resource allocation in channel/router (flow-control digit)
- Phit = Smallest unit of data processed by a channel/router (physical digit)

For the processor-memory networks we will primarily study:

$$\text{Message} = \text{Packet} = \text{Flit}$$

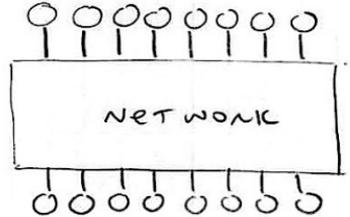
The processor will send memory request messages/packets over the network, and the memories will send memory response messages/packets back to the processor. In this context, we consider the network messages/packets to be the “transactions”.

Processors	:	Instructions
Memories	:	Memory accesses
Networks	:	Network packets

2. Network Topology

Network topology is the arrangement of channels and routers to interconnect sources and destinations. We will explore four different topology classes:

- Single-stage bus topologies
- Single-stage crossbar topologies
- Multi-stage butterfly topologies
- Multi-stage torus topologies



2.1. Single-Stage Bus Topology

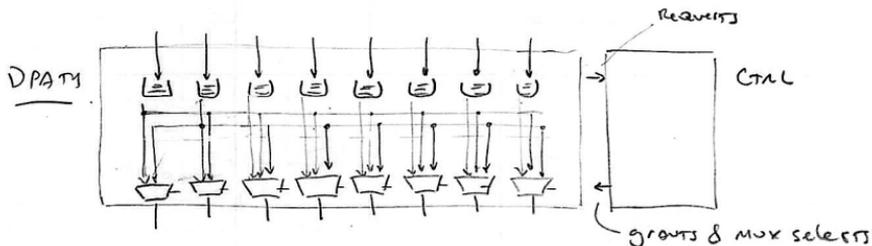
SWMR Bus

MWSR Bus

MWMR Bus

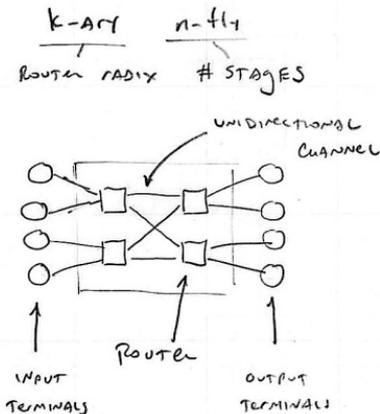
Integrating Multiple Buses

2.2. Single-Stage Crossbar Topology



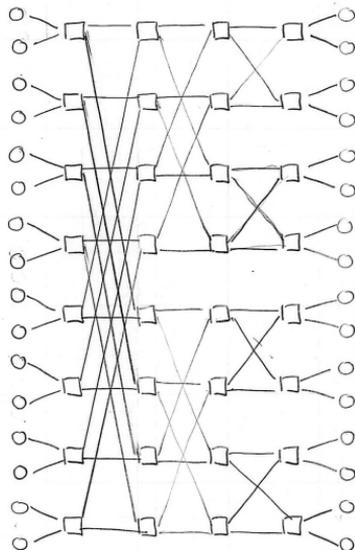
- Single-stage topologies are difficult to scale in terms of cycle time, energy, and area
- Multi-stage topologies improve scalability but raise many other interesting challenges

2.3. Multi-Stage Butterfly Topology



2-ARY 2-FLY

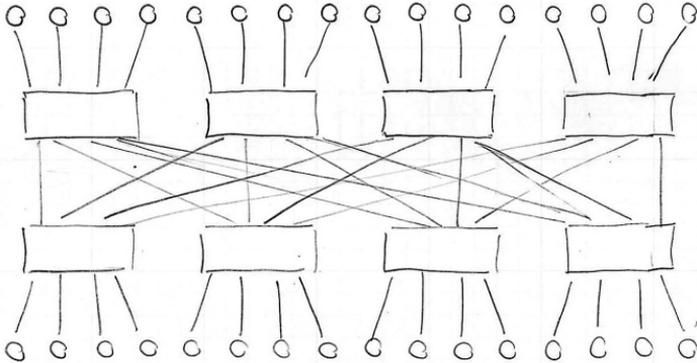
EACH ROUTER IS SIMILAR TO A 2x2 CROSSBAR



2-ARY

4-FLY

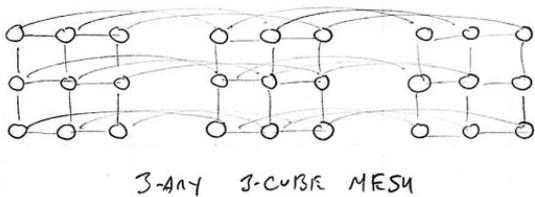
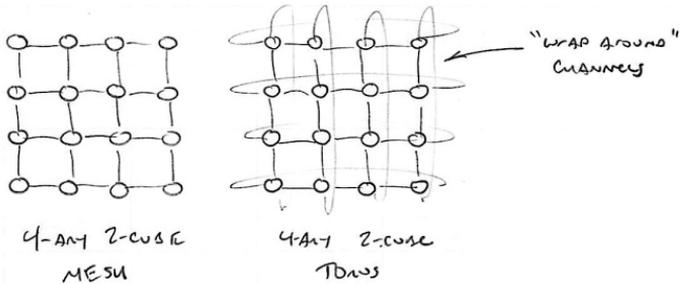
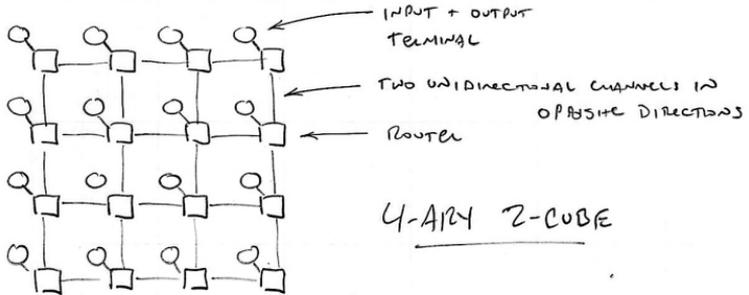
$2^4 = \text{NODES}$

Example Butterfly Topology: 4-ary 2-fly**Example Butterfly Topology: 3-ary 2-fly**

Sketch a 3-ary 2-fly. Use circles for the terminals, squares for the routers, and lines for one uni-directional channel.

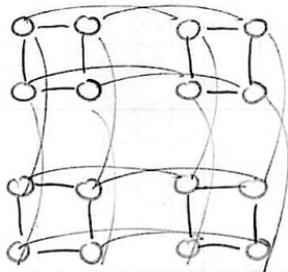
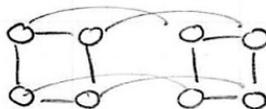
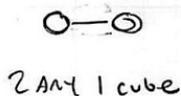
2.4. Multi-Stage Torus Topology

K-ARY N-CUBE
 / \
 NODES IN EACH N DIM GRID
 DIMENSIONS



Example Torus Topology: 8-ary 1-cube torus

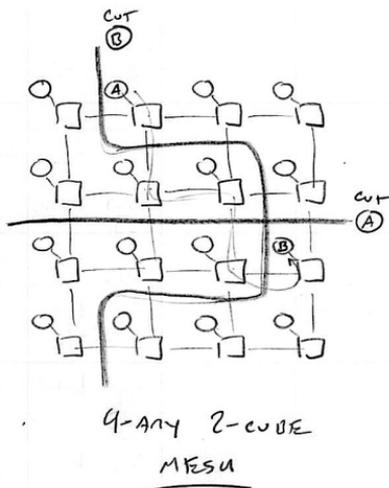
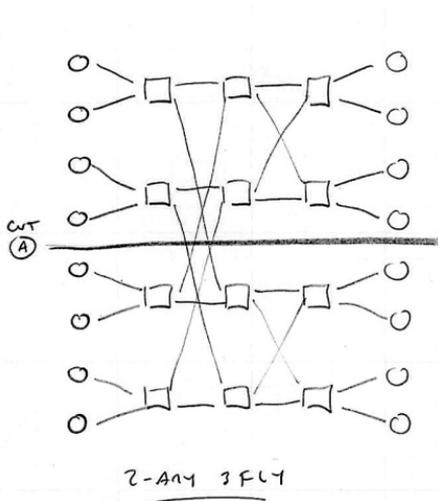
Sketch a 8-ary 1-cube torus. Use circles for a terminal+router and lines for the two uni-directional channels in opposite directions.

Constructing k-ary n-cube from k k-ary (n-1) cubes

HYPER CUBES

Binary 1-cube

2.5. Terminology



Nodes and Channels

- Uni-directional channels in butterfly
- Bi-directional channels in mesh
- Indirect Network: node is either a terminal or router (butterfly)
- Direct Network: node combines a terminal and router (mesh)
- Channel parameters
 - w_c = channel width (number of pins/wires)
 - f_c = channel frequency
 - t_c = channel latency
 - l_c = channel length
 - b_c = channel bandwidth ($w_c \times f_c$)

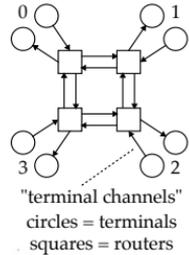
Bisection Cuts

- Cut: set of channels that partitions terminals into two sets
- Bisection Cut: cut that partitions terminals in two equal halves
- Min Bisection Cut (B_c): min channel count over all bisection cuts
- Bisection Bandwidth (B_B): min bandwidth over all bisection cuts
- For networks with uniform channel bandwidth: $B_B = B_c \times b_c$
- Bisection bandwidth is a good way to estimate global wiring resources (i.e., technology constraints)
- Example from previous butterfly/mesh topologies
 - Cut A on butterfly is a minimum bisection cut
 - Both cut A and B on mesh are bisection cuts
 - Cut A on mesh is a minimum bisection cut

Paths

- Channel Hop Count (H_c): number of channels on a path
- Channels from terminal to first router may or may not be included
- Router Hop Count (H_r): number of routers on a path
- Minimal Path: smallest hop count between two terminals
- Diameter (H_{max}): largest minimal path between all terminal pairs
- Average min hop count (H_{min}): average over all terminal pairs
- Example from previous butterfly/mesh topologies
 - Path from A to B on mesh: $H_r = 5, H_c = 4$
 - $H_{max,bfly} = 4, H_{max,mesh} = 8$
 - $H_{min,bfly} = 4$

Calculating H_{min} usually requires enumerating the minimal hop count for every source/destination pair. The book states that H_{min} for a torus with even k is $nk/4$ and for a mesh with even k is $nk/3$. Where does this come from?



Calculating H_{min}

	src						
	0	1	2	3			
DEST 0	1	2	3	2	$1 \times 4 = 4$	including $i \rightarrow i$	$32/16 = 2$
1	2	1	2	3	$2 \times 8 = 16$		
2	3	2	1	2	$3 \times 4 = 12$	excluding $i \rightarrow i$	$28/12 = 2.3$
3	2	3	2	1			

Calculating $H_{min,c}$ with terminal channels

	src						
	0	1	2	3			
DEST 0	2	3	4	3	$2 \times 4 = 8$	including $i \rightarrow i$	$40/16 = 2.5$
1	3	2	3	4	$3 \times 8 = 24$		
2	4	3	2	3	$4 \times 4 = 16$	excluding $i \rightarrow i$	$40/12 = 3.3$
3	3	4	3	2			

Calculating $H_{min,c}$ without terminal channels

	src						
	0	1	2	3			
DEST 0	0	1	2	1	$1 \times 8 = 8$	including $i \rightarrow i$	$16/16 = 1$
1	1	0	1	2	$2 \times 4 = 8$	excluding $i \rightarrow i$	$16/12 = 1.3 <$
2	2	1	0	1			
3	1	2	1	0			

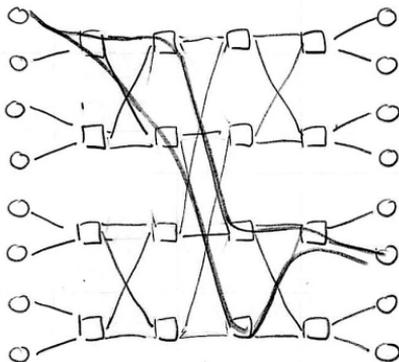
EQUATION FROM BOOK IS $\frac{nk}{3}$ FOR EVEN n

$$\frac{nk}{3} = \frac{2 \cdot 2}{3} = 1.33$$

THIS IS FOR $H_{min,c}$
WITHOUT terminal channels
+ ignoring i sending to itself

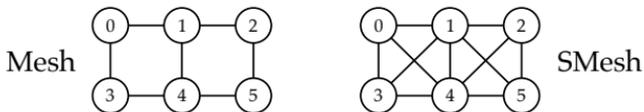
Path Diversity

- Path diversity is number of paths between any two terminals
- Mesh has high path diversity, butterfly has no path diversity



ADDING EXTRA
BFLY STAGES
↑ PATH DIVERSITY

Example Topologies



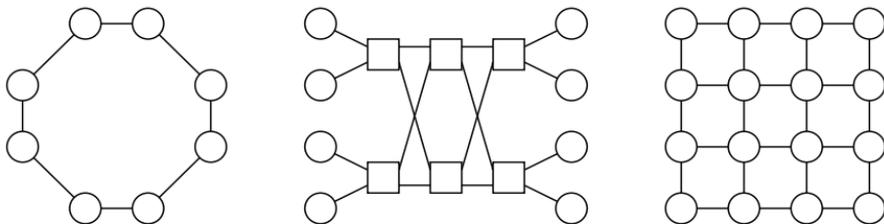
Assume b_c for mesh is 32 bits/cycle, while b_c for smesh is 16 bits/cycle. smesh has reduced channel bandwidth due to increased number of channels, assumes specific technology constraint. Calculate the following parameters for each topology.

	B_C	B_B	H_{max}	H_{min}
mesh				
smesh				

3. Network Routing

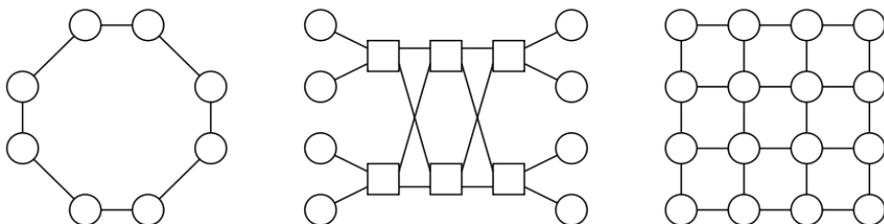
Network routing is the choice of path from source to destination along channels and routers. Routing algorithms can be oblivious (not factor in network state) or adaptive (factor in network state), and can also be deterministic or non-deterministic.

3.1. Oblivious Deterministic Routing



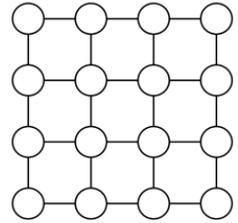
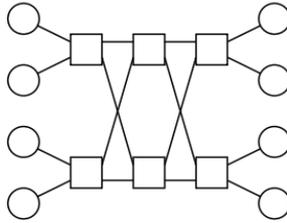
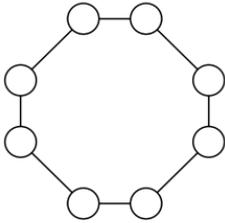
- Ring: always use minimal path, equidistant choose CW
- Butterfly: use destination to choose middle router
- Mesh: route in X first, then route in Y (Dimension-Ordered-Routing)

3.2. Oblivious Non-Deterministic Routing



- Ring: randomly choose CW vs CCW
- Butterfly: randomly choose middle router
- Mesh: randomly choose between XY-DOR and YX-DOR (O1-TURN)

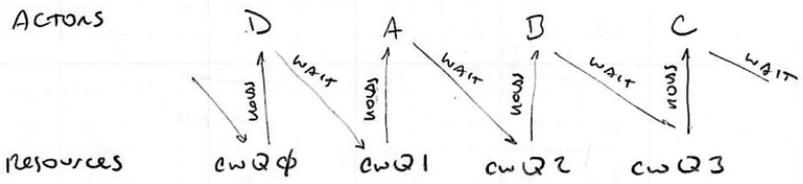
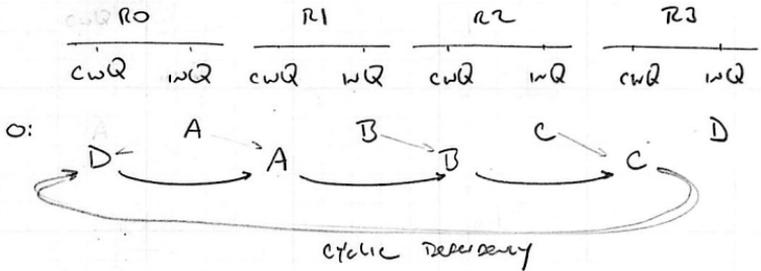
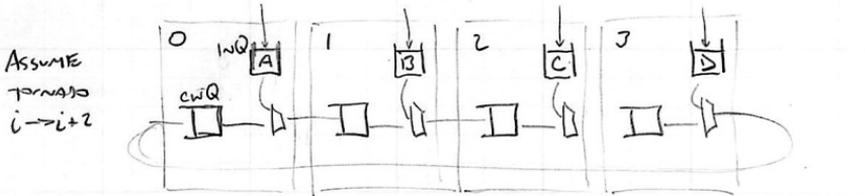
3.3. Adaptive Routing



- Ring: look at queues, choose direction with least congestion
- Butterfly: look at queues, choose middle router with least congestion
- Mesh: look at queues during each hop to choose X vs Y

3.4. Deadlock

Cyclic Dependency in "Waits For" AND "Holds" Relation



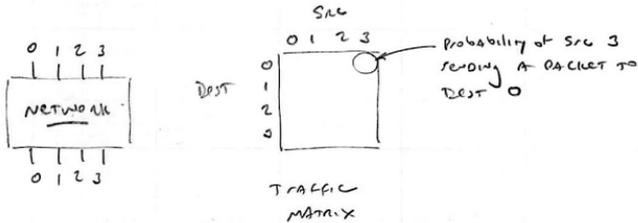
- DEADLOCK AVOIDANCE vs DEADLOCK DETECTION / RECOVERY

CONTRAST TO LIVELOCK
 CHAOS ROUTING ALGORITHM

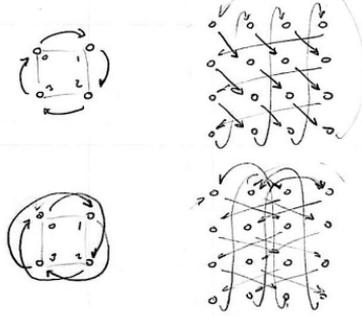
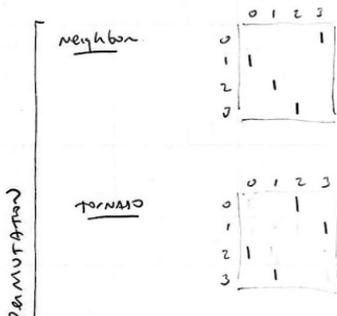
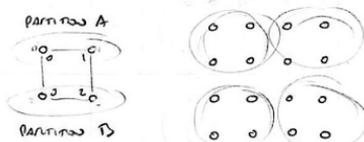
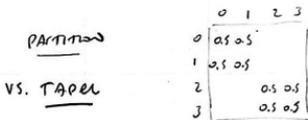
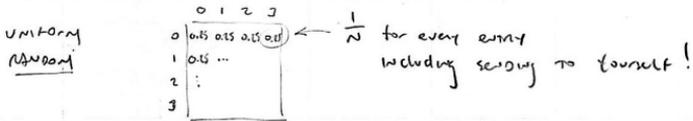
4. Analyzing Network Performance

Similar to analyzing processors and memories, we will use simple first-order equations to help build intuition about throughput and latency trade-offs in networks.

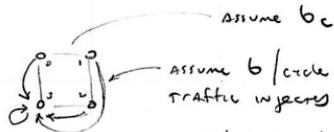
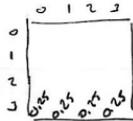
4.1. Traffic Patterns



Admissible Traffic Patterns



Hot Spot

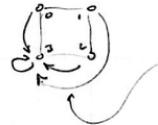
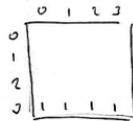


assume b_c
assume b / circle traffic injects

$0.25b + 0.25b = 0.5b$
going over $2 \rightarrow 3$ channel

INADMISSIBLE TRAFFIC PATTERNS

oversubscribes
Hotspot

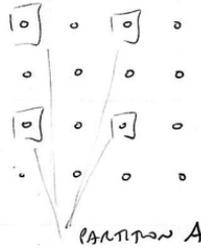
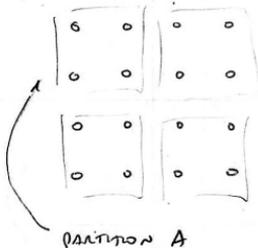


$1b + 1b = 2b$
going over the $2 \rightarrow 3$ channel

Channel $2 \rightarrow 3$ oversubscribed
but so is ejected channel
to output terminal at node 3

LOGICAL TO PHYSICAL MAPPING

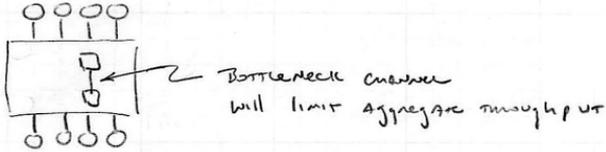
ASSUME PARTITIONED TRAFFIC PATTERNS
How are logical src/dest IDs in traffic patterns
MAPPED TO PHYSICAL TERMINAL IDS?



MAPPING CAN TURN ANY LOGICAL PERMUTATION PATTERN INTO ANY OTHER PATTERN.

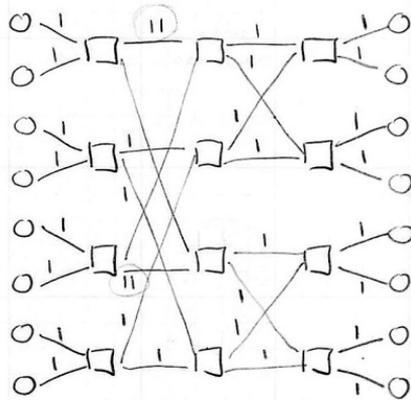
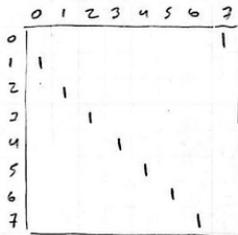
gives PERMUTATION PATTERN usually assumes a specific MAPPING

4.2. Ideal Throughput



CHANNEL LOAD (λ_c) IS AMOUNT OF TRAFFIC THAT CROSSES CHANNEL c IF EACH INPUT INJECTS ONE UNIT OF TRAFFIC ACCORDING TO GIVEN TRAFFIC PATTERN

Rx Traffic Pattern
Src $i \rightarrow$ Dest $i+1$



CHANNEL LOAD RANGES FROM 0-2

MAX CHANNEL LOAD (λ_{\max}) IS 2. THESE ARE THE BOTTLENECK CHANNELS THAT WILL LIMIT THROUGHPUT

ALTERNATIVE WAY TO THINK ABOUT λ .

CHANNEL LOAD IS RATIO OF BW DEMANDED FROM CHANNEL TO INPUT BW INJECTED BY ONE TERMINAL

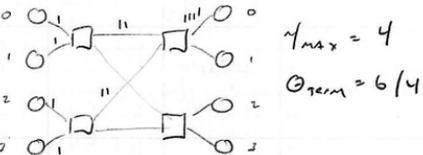
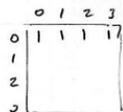
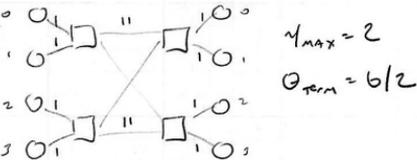
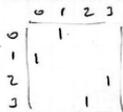
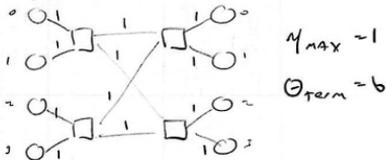
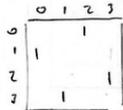
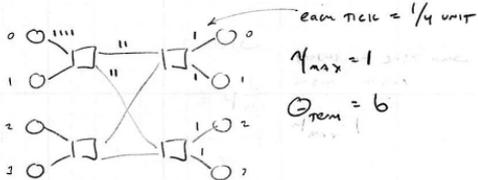
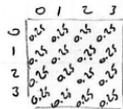
IDEAL THROUGHPUT $\Theta_{\text{term}} = \frac{b_c}{\eta_{\text{max}}}$

\leftarrow Bw of bottleneck channel
 \leftarrow max customer load

in previous example

$\Theta_{\text{term}} = \frac{6}{\eta_{\text{max}}} = \frac{6}{2}$ $\Theta_{\text{tot}} = N \cdot \frac{6}{\eta_{\text{max}}} = 8 \cdot \frac{6}{2} = 46$

SMALL DFLY EXAMPLES



More generally for uniform random traffic

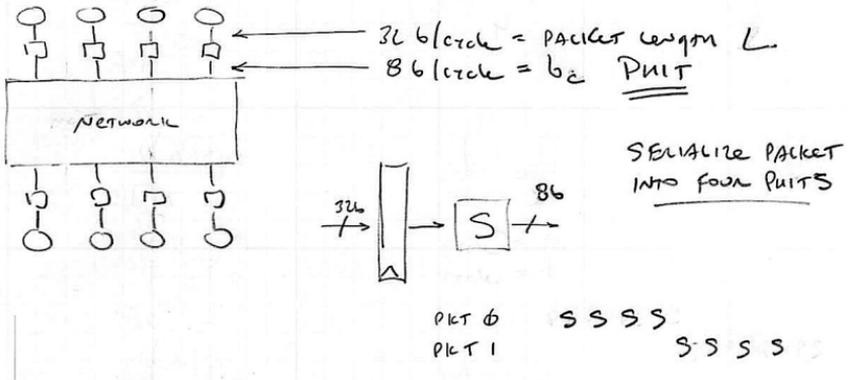
- on average with uniform random traffic, HALF the traffic crosses the bottleneck
- N TOTAL UNITS of traffic, $N/2$ cross bottleneck
- ideal routing will evenly BALANCE load across bottleneck channels (THIS WAS AN ISSUE in Ring example)

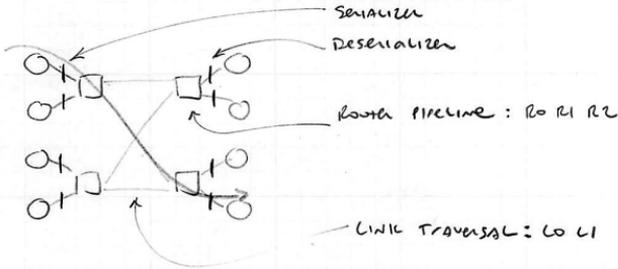
$$\eta_{\max} = \frac{N/2}{B_c} = \frac{N}{2B_c}$$

$$\Theta_{\text{form}} = \frac{6}{N/2B_c} = \frac{26B_c}{N} = \frac{2B_B}{N}$$

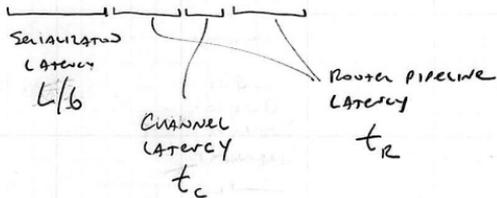
$$\Theta_{\text{TOT}} = N \frac{2B_B}{N} = 2B_B$$

4.3. Zero-Load Latency





PKT | HEAD PUIT → S R0 M R2 L0 L1 R0 R1 R2 D
 BODY PUIT S R0 M R2 L0 L1 R0 R1 R2 D
 BODY PUIT S R0 M R2 L0 L1 R0 R1 R2 D
 TAIL PUIT S R0 M R2 L0 L1 R0 R1 R2 D →



$$T = T_{\text{HEAD}} + \frac{L}{b}$$

T_{HEAD} → serialisation Latency
 ↳ HEAD PUIT Latency includes t_c, t_r , hop count, + contention

$$T_{\phi} = H_R t_r + H_C t_c + \frac{L}{b}$$

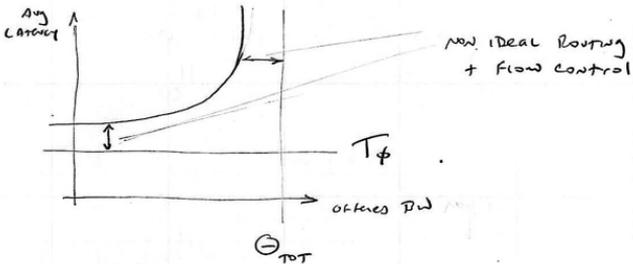
$H_R t_r$ → Latency due to Router hops
 $H_C t_c$ → Latency due to channel hops
 $\frac{L}{b}$ → Serialisation Latency
Zero load latency (no contention)

FOUR WAYS TO IMPROVE LATENCY

$$T_{\phi} = H_n t_n + H_c t_c + \frac{L}{b}$$

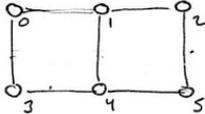
Shorter routes faster routes faster channels with overhead on shorter msg

Avg LATENCY vs OFFERED TRW

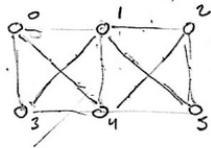


4.4. Comparing Topologies

MESH



SMESH



ASSUME:

$$L = 1286$$

b_c for mesh is 32 b/cycle

b_c for smesh is 16 b/cycle

$$t_n = 3$$

$$t_c = 1$$

← reduces channel bandwidth due to increases # channels
Assuming constant resources

WHICH TOPOLOGY CAN ACHIEVE HIGHER IDEAL THROUGHPUT UNDER UNIFORM RANDOM TRAFFIC?

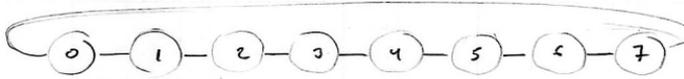
$$(a) \quad \Theta_{\text{term, mesh}} > \Theta_{\text{term, smesh}}$$

$$(b) \quad \Theta_{\text{term, mesh}} < \Theta_{\text{term, smesh}}$$

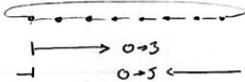
$$(c) \quad \Theta_{\text{term, mesh}} = \Theta_{\text{term, smesh}}$$

CALCULATE ZERO LOAD LATENCY UNDER UNIFORM RANDOM TRAFFIC

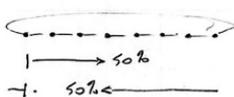
4.5. Comparing Routing Algorithms



- GREEDY: Always use minimal path, equivalent choose randomly

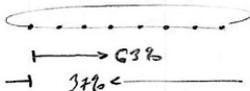


- UNIFORM RANDOM: RANDOMLY PICK DIRECTION



for $0 \rightarrow 3$
 50% TAKE 4 HOPS
 50% TAKE 6 HOPS

- WEIGHTED RANDOM: RANDOMLY PICK DIRECTION BUT WEIGHT PROBABILITY BY DISTANCE



for $0 \rightarrow 3$
 $5/8$ TIME TAKE 3 HOP PATH
 $3/8$ TIME TAKE 5 HOP PATH

PROBABILITY OF TAKING SHORTEST PATH IS

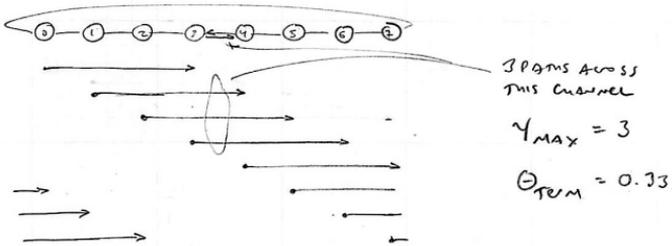
$$\frac{N - H_{\min, R} + 1}{N}$$

- ADAPTIVE: LOOK AT QUEUES IN EITHER DIRECTION SEND IN DIRECTION OF QUEUE THAT HAS MOST FREE SLOTS.

DO NOT CHANGE DIRECTION AFTER INITIAL CHOICE

- Evaluate tornado traffic pattern on the 8-node ring
- Recall that in tornado, node i sends to $i + ((N - 1)/2) \bmod N$

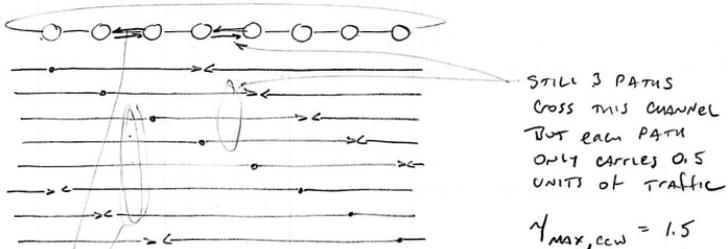
greedy Algo



Clockwise channels NOT USED AT ALL!
POOR LOAD BALANCING

$$T_0 = 4$$

RANDOM ALGO



WHAT ABOUT CLOCKWISE CHANNELS THOUGH?

5 PATHS CROSS EACH CW CHANNEL
EACH IS 0.5 UNIT OF TRAFFIC

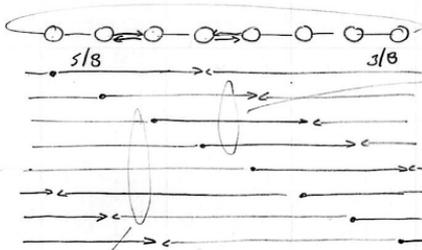
$$Y_{\max,cw} = 2.5 > Y_{\max,cw} = 1.5 \quad \Theta_{\text{TORN}} = 0.4$$

NOW THE CLOCKWISE CHANNELS ARE THE BOTTLENECK!

$$T_{\Phi} = 0.5 \times 4 + 0.5 \times 6 = 5$$

Weighted Round Algo

SAME NUMBER OF PATHS CROSS CW AND CCW DIRECTION AS BEFORE EXCEPT NOW WITH DIFFERENT AMOUNTS OF TRAFFIC PER PATH.



3 PATHS EACH WITH $5/8$ TRAFFIC

$$N_{\max, \text{ccw}} = 3 \cdot \frac{5}{8} = 1.875$$

5 PATHS EACH WITH $3/8$ TRAFFIC

$$N_{\max, \text{cw}} = 5 \cdot \frac{3}{8} = 1.875$$

$$N_{\max, \text{cw}} = N_{\max, \text{ccw}}$$

BALANCED!

$$G_{\text{rem}} = \frac{1}{1.875} = 0.53$$

$$T_0 = \frac{5}{8} \times 4 + \frac{3}{8} \times 6 = 2.5 + 2.25 = 4.75$$

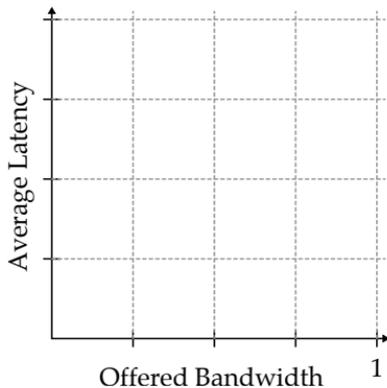
ADAPTIVE

MINIMAL LATENCY AT LIGHT LOAD
MAX THROUGHPUT AT HIGH LOAD

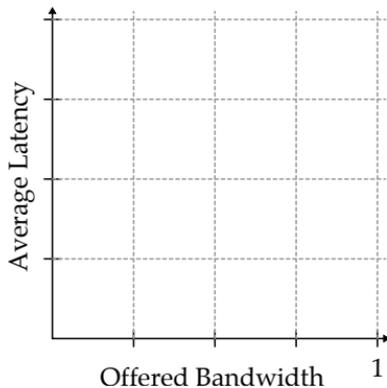
Let's fill in the following table to summarize the performance of the four routing algorithms on the 8-node given the tornado and uniform random traffic patterns.

Routing Algo	Tornado		Uniform Random	
	Θ_{term}	T_0	Θ_{term}	T_0
Greedy			1.00	3.0
Uniform Random			0.57	4.5
Weighted Random			0.76	3.1
Adaptive			1.00	3.0

Tornado Traffic Pattern
Bandwidth vs. Latency



Uniform Random Traffic Pattern
Bandwidth vs. Latency



Activity: Routing on butterfly with extra stage

Consider two routing algorithms for a 2-ary 2-fly with an extra stage:

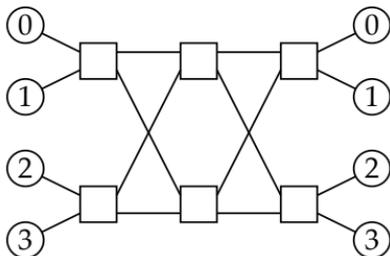
- Destination Tag Routing: use destination to choose middle router
- Random Middle Router: randomly choose middle router

Compare these two routing algorithms in terms of ideal terminal throughput (Θ_{term}) and zero-load latency (T_0) for the following permutation traffic pattern.

- src 0 \rightarrow dest 1
- src 1 \rightarrow dest 0
- src 2 \rightarrow dest 3
- src 3 \rightarrow dest 2

Hint: Use the tick-mark method to calculate the max channel load for each routing algorithm.

Destination-Tag Routing



Random Middle Routing

